Exhibit 16

Moving From Detection to Prevention of Modern Malware

WildFire: Modern Malware Prevention

November 2012

Palo Alto Networks
3300 Olcott St
Santa Clara, CA
95054
www.paloaltonetworks.com

Table of Contents

Modern Malware: The Heart of Network Attacks

Modern malware is at the center of today's most sophisticated attacks, enabling attackers to gain a foothold and persist within an enterprise which they can use to dig deeper into the network, control their attack, and steal information over a period of weeks, months, or even years. Furthermore, this malware puts the attacker inside a trusted position in the network, allowing for a sophisticated network attack that is performed from the inside out. Malware also has the advantage of being extremely easy to modify and thus avoid security controls. Malware can be re-encoded, recompiled or further customized for the target in order to avoid antivirus signatures. This combination makes malware both one of the most powerful parts of an attack, but also one of the most difficult to detect.

As a result, IT security teams simply must be able to address these new threats proactively as a part of their everyday security practice, or risk ceding the intrusion prevention battle to the attackers. This requires new solutions and techniques designed to control these modern threats, but also meet the established enterprise requirements of accuracy, prevention and scalability.

Why Existing Solutions Can't Fix The Problem

Traditional network security products have a variety of deficiencies that limit the ability to control modern and evolving malware. Traditional firewalls only look at packet headers, and generally were not designed to detect malware. Other network-based anti-malware solutions are proxy-based, which are typically prohibitively slow and limited to only protecting against web-based malware. However, antivirus web-proxies, IPS products and email gateway products share a far more serious and fundamental flaw – namely they only protect against known malware that has been submitted to the vendor or captured in a honeypot. These generic detections often have little relevance to the new, polymorphic or targeted malware that hits a given network on a daily basis. In short, these solutions only protect against things that are already known to be bad, and lack the ability to detect new threats hitting the network.

Another key challenge in the control of modern malware is the seemingly simple requirement of visibility into traffic. As an example, as much as 33% of enterprise network traffic is protected by SSL encryption, which can prevent security solutions from inspecting the traffic within. This trend is expected to continue, as many web-based email, social networking, and other Enterprise 2.0 web applications default to HTTPS to protect data in transit. While this technology obviously provides benefits in the form improved session privacy for the connection, it also provides an encrypted channel to distribute malware to hosts on a network. This is just the tip of the iceberg, as malware and their authors have increasingly adopted a variety of additional techniques to obscure both the infecting files as well as the ongoing command and control traffic that modern malware depends on. This includes tunneling communications within approved protocols or traffic as well as using proxies, circumventors and non-standard ports in order to avoid traditional security solutions. These techniques, like SSL allow threats to remain hidden even as they repeatedly cross the perimeter without inspection.

While more traditional solutions certainly have their own set of problems, many newer solutions with more advanced analysis capabilities fail to meet the standard of quality needed in enterprise deployments in terms of the ability to analyze all traffic and threats, actually prevent detected threats, and scale to meet the traffic demands of a modern enterprise network. To meet the challenge a new solution is required that combines new automated malware analysis techniques with industry-standard levels of prevention, performance and scalability.

Palo Alto Networks WildFire®

To meet the challenge of modern-day malware, Palo Alto Networks has developed WildFire, which extends the power of the next-generation firewall to detect and prevent modern malware by directly executing them in a virtual environment and observing malicious behavior. IT teams are notified within minutes of analysis, and true malware protections are created, tested and delivered to all WildFire subscribers within an hour of initial detection. This enables Palo Alto Networks to identify and control malware quickly and accurately, even if the malware has never been seen in the wild before.

WildFire combines the unmatched visibility and enforcement of the next-generation firewall with the unlimited scalability of the cloud to deliver a solution to modern malware that is reliable, practical and actively protects the network. WildFire begins by leveraging the power of a customer's on-premises firewalls to ensure full visibility of all traffic to detect malware on non-standard ports, hidden with tunnels or SSL encryption. Unknown files are securely forwarded to the WildFire cloud where they can be safely executed in a virtual environment, and where true malware signatures are developed and regression tested. Protections are delivered back to all firewalls, worldwide where enforcement is again performed by the firewall at speeds up to 10Gbps.

- **Deep Visibility:** The WildFire solution makes extensive use of Palo Alto Networks App-ID analyze all traffic and detect threats even on non-standard ports. WildFire supports a variety of applications including email, web, FTP and SMB. Additionally, on-device SSL decryption enables administrators to configure policies that detect file transfers through HTTPS-encrypted web applications and send them to WildFire for analysis.

- **Virtualized Sandbox:** When the firewall encounters an unknown file, the file can be submitted to the WildFire virtualized sandbox, which is continually maintained by Palo Alto Networks threat researchers. Submissions can be made manually or automatically based on policy. Each sample is executed in a virtual machine where Palo Alto Networks can directly observe more than 100 malicious behaviors that can reveal the presence of malware. As malware behaviors change and develop new anti-analysis techniques, Palo Alto Networks can update analysis logic and technology to keep pace with malware without requiring any updates to end-user hardware.

- **Automated Signature Generator:** When a sample is identified as malware, it is passed on to a signature generator, which automatically generates a signature for the sample and tests it for accuracy. Signatures are based on a combination of header and body content of the malicious file, enabling a single signature to address multiple variants of a malware sample and to resist common repackaging and polymorphism techniques. With WildFire in the cloud, signatures are automatically regression tested against an extensive database of samples, and then delivered to all WildFire subscribers worldwide within an hour of the initial detection. Palo Alto Networks also generates signatures for the all-important command and control traffic, malicious DNS queries and any URLs used by the malware allowing staff to disrupt and quarantine any existing infections.

- **True In-line Enforcement:** When protections are delivered to customer firewalls, enforcement is handled by the proven and 3rd party tested threat prevention engine. Malware enforcement is handled by the stream-based malware engine, which enables the in-line blocking of malware without the negative performance impacts commonly associated with proxy-based solutions. As the in-line firewall, Palo Alto Networks reserves the ability to actually drop network traffic instead of relying only on TCP resets, which can be easily filtered or ignored by malicious end-points.

- **Actionable Intelligence:** In addition to protection, administrators have access to a wealth of actionable information about the detected malware through the WildFire portal. A detailed behavioral report of the malware is produced, along with information on the user that was targeted, the application that delivered the malware, and all URLs involved in the delivery or phone-home of the malware. Additionally WildFire subscribers receive integrated WildFire logs within minutes of the malware being analyzed, allowing staff to instantly see results and correlate WildFire findings with application, URL, users, file and threat logs.

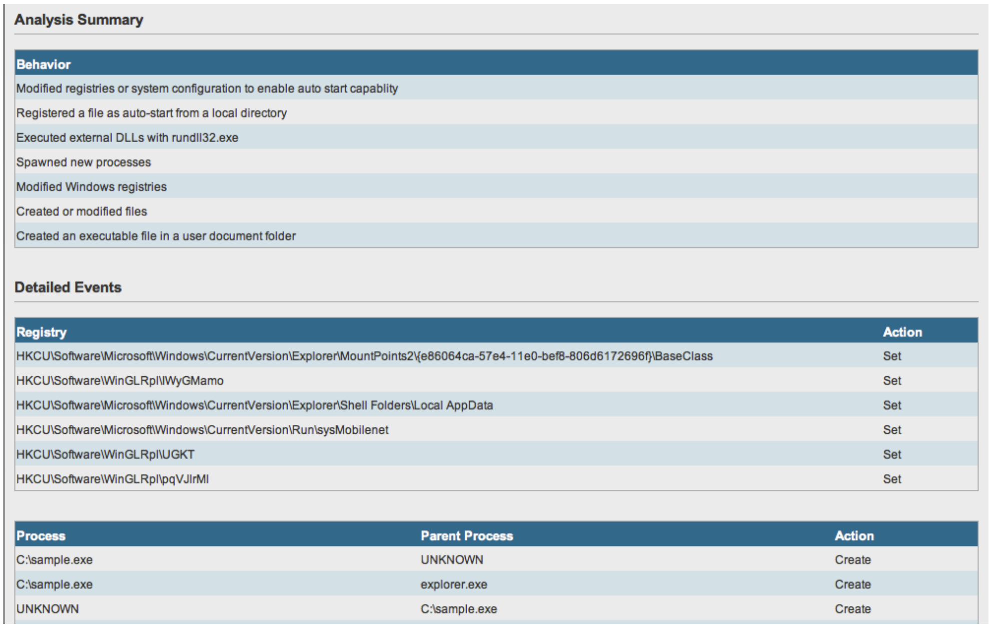

Analysis Summary

Behavior
Modified registries or system configuration to enable auto start capablity
Registered a file as auto-start from a local directory
Executed external DLLs with rundll32.exe
Spawned new processes
Modified Windows registries
Created or modified files
Created an executable file in a user document folder

Detailed Events

Registry	Action
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\MountPoints2\{e86064ca-57e4-11e0-bef8-806d6172696f}\BaseClass	Set
HKCU\Software\WinGLRpl\IWyGMamo	Set
HKCU\Software\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders\Local AppData	Set
HKCU\Software\Microsoft\Windows\CurrentVersion\Run\sysMobilenet	Set
HKCU\Software\WinGLRpl\UGKT	Set
HKCU\Software\WinGLRpl\pqVJIrMI	Set

Process	Parent Process	Action
C:\sample.exe	UNKNOWN	Create
C:\sample.exe	explorer.exe	Create
UNKNOWN	C:\sample.exe	Create

Excerpt from a malware analysis report

- **Behavioral Botnet Report:** In addition to the direct analysis of malware in WildFire, the Palo Alto Networks solution also includes the ability to identify the presence of modern malware through the monitoring and correlation of suspicious network traffic. The behavioral botnet report looks for a variety of telltale signs of a botnet infection, such as the presence of unknown application traffic, IRC traffic, repeated attempts to download files, and connections to unknown or newly registered domains. The report leverages User-ID to identify the infected user and the factors that contributed to the analysis.

WildFire In Action

WildFire is easily put into action by configuring a simple policy on a Palo Alto Networks next-generation firewall. Firewall policies govern exactly what types of files are submitted for analysis and any correlating information that should be included or not. When the firewall discovers an unknown file, the firewall can query the cloud to see if the file's hash has been analyzed previously, avoiding the potential to reanalyze files repeatedly. If the file has not been analyzed, the file is securely uploaded to the WildFire cloud via connection secured by certificates on both sides that are signed by Palo Alto Networks to prevent the potential for a man-in-the-middle event.

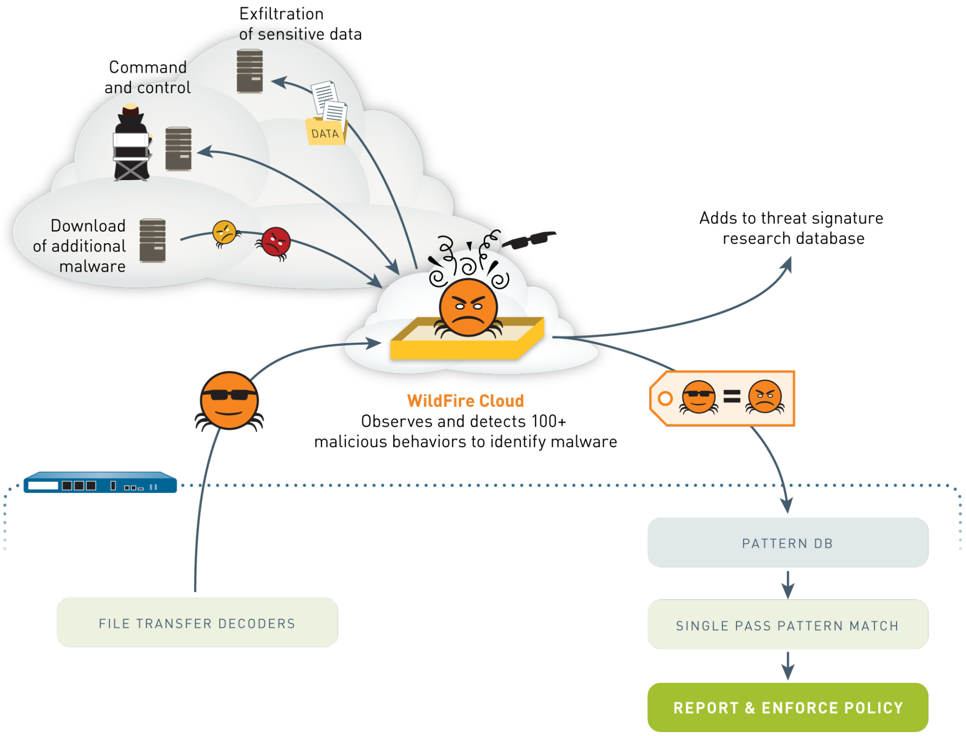

When a new sample is sent to WildFire, the file is executed within a virtual machine, and analyzed for malicious behavior. WildFire monitors activity on the virtual machine, looking for over 100 malicious behaviors that can indicate the presence of malware, including modifications to the host system, injection of code into other processes, evasion attempts, attempts to subvert local security controls as well as a variety of malicious network and hacking activity. When the analysis is complete, WildFire determines whether the sample is benign or malicious, and a log of the analysis is delivered to the WildFire subscriber within minutes of the initial submission. IT teams can also elect to have notifications delivered via email based on policy.

WildFire automatically generates signatures for samples that are malicious, which are immediately regression tested against an extensive database of clean and malicious samples to ensure signature quality. These malware signatures are based a combination of unique identifiers in the file header and body to ensure the signature will apply to samples regardless of name changes, byte-padding and basic polymorphic techniques. These signatures are delivered to all WildFire subscribers within 1 hour of the initial submission of the sample to WildFire, allowing other WildFire customers to be protected from rapidly spreading malware. In addition to creating signatures for the infecting files, Palo Alto Networks also creates signatures that identify the command and control traffic of the malware, ensuring that staff can also instantly stop and quarantine any active threats already in the network. The signatures of all malicious WildFire submissions are then bundled with the daily antivirus updates and distributed to Palo Alto Networks customers that have a threat protection subscription.

Advantages of Combating Malware in the Cloud

Other virtualized malware detection products have introduced substantial hardware, financial and management burden by requiring dedicated hardware appliances at every ingress point to be protected. WildFire, on the other hand, leverages the reach and scalability of the cloud while remaining tightly integrated with the Palo Alto Networks next-generation firewalls for simple setup and enforcement. All that is required to protect against zero-day malware is a Palo Alto Networks next-generation firewall.

The virtualized analysis of malware, needless to say makes heavy use of virtual systems where malware is actively analyzed. Cloud computing has become a dominant force in computing precisely for its ability to dynamically scale and add capacity without being limited by local, on-premise hardware. Alternate architectures attempt to leverage virtualization for malware analysis, but are hard limited by the hardware on-site and are quickly overwhelmed by real-world traffic flows and malware analysis requirements. In short a local hardware-based architecture creates a very expensive chokepoint for malware analysis that actually decreases accuracy, whereas the WildFire combination of firewall and cloud maintains high performance and unlimited analysis capacity.

Additionally, WildFire can be easily updated by Palo Alto Networks researchers to quickly respond to evolving malware strategies, removing the need for IT teams to constantly update software or service packs on an in-house sandbox. Malware techniques evolve rapidly, and in many cases new techniques are directed at avoiding virtualized analysis. The WildFire cloud offloads these complexities to Palo Alto Networks experts who continually update the power of WildFire while requiring no updates to the end-user firewall.

Offloading the complex infrastructure and compute resources required to host advanced virtualized based malware detection is not the only advantage. With WildFire in the cloud, Palo Alto Networks breaks the silos of information that have traditionally plagued other attempts at malware detection. In short, if a new or targeted threat is detected, that information and the ability to protect against the threat needs to be shared across the entire enterprise and not limited to the ingress point that detected the threat. WildFire centralizes the analysis of unknown files, and also provides a centralized source of protections for all Palo Alto Networks firewalls. Users of WildFire automatically receive signatures for never-before-seen malware that other users in the network have submitted to WildFire as part of their threat prevention subscription service.

Fully Integrated Threat Prevention

For all of the unique power and visibility provided by WildFire, it is important to remember that it only represents part of the threat prevention puzzle. The control of unknown malware must be integrated into an overall security strategy that incorporates application controls, known exploits and malware, dangerous URLs and websites, and even dangerous file types or restricted content. The Palo Alto Networks next-generation firewall integrates all of these technologies into a unified view and approach to policy. Furthermore, each discipline provides value to the others. When WildFire detects unknown malware, the firewall allows users to see the application that delivered the file, the user that was targeted, the website that delivered the file and the URLs that it attempted to contact. This provides the context that is needed in order to deliver a coordinated and comprehensive response that goes beyond stopping individual threats and instead controls the overall risk of the enterprise.

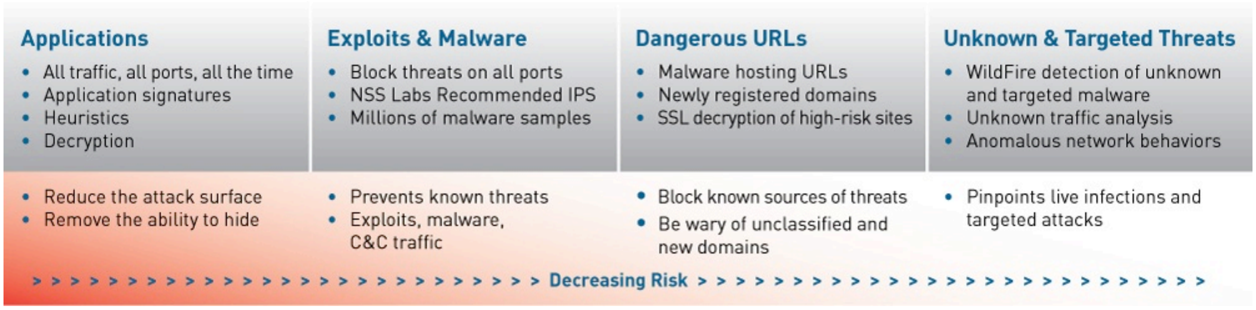

Applications	Exploits & Malware	Dangerous URLs	Unknown & Targeted Threats
• All traffic, all ports, all the time • Application signatures • Heuristics • Decryption	• Block threats on all ports • NSS Labs Recommended IPS • Millions of malware samples	• Malware hosting URLs • Newly registered domains • SSL decryption of high-risk sites	• WildFire detection of unknown and targeted malware • Unknown traffic analysis • Anomalous network behaviors
• Reduce the attack surface • Remove the ability to hide	• Prevents known threats • Exploits, malware, C&C traffic	• Block known sources of threats • Be wary of unclassified and new domains	• Pinpoints live infections and targeted attacks

> Decreasing Risk >

In the end, the detection and prevention of new and evolving malware must be integrated into your existing threat prevention and security process. Each discipline should be reliable, fast, provide unique value, and correlate with the other available techniques. This ensures the rollout of security technology that not only addresses the threat, but also supports your security process.

Exhibit 17

paloalto networks®

Palo Alto Networks® Threat Visibility for Government Networks

Learning About Your Risks

Discovering an Unknown Application or File Type

Having established your application whitelisting and created policies for the applications and file types allowed on each segment of your government network, you are in the best position to identify any anomalous applications and files. The Palo Alto Networks enterprise security platform prevents all known file types from traversing your network in-line. Furthermore, all unknown files are dynamically analyzed by the Wildfire® Threat Intelligence Cloud, where unknown threats are identified, and protections are created and shared with all customers within as little as 30 minutes.

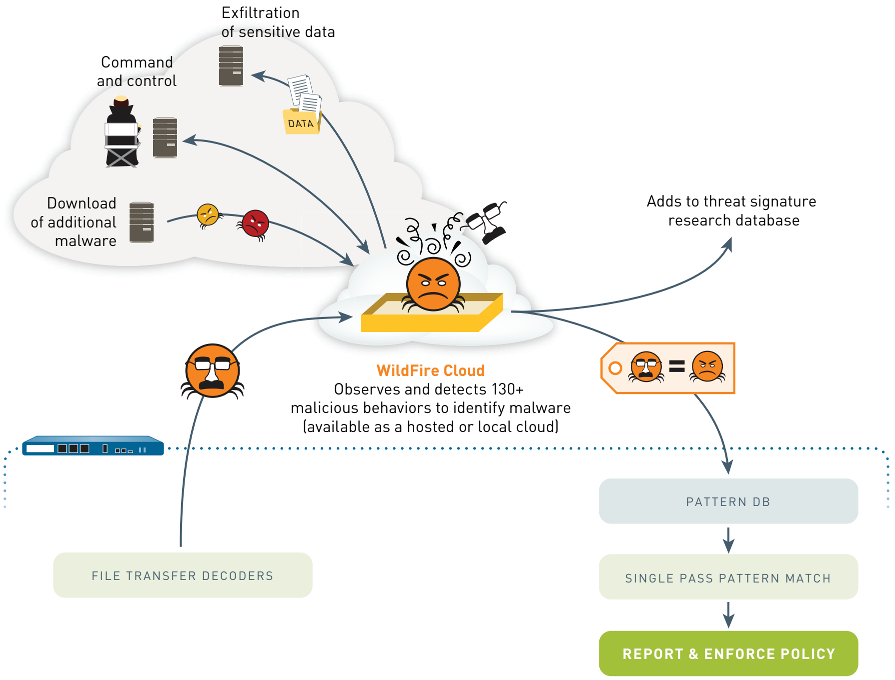

The Palo Alto Networks single pass software checks your policies for the allowed applications, content, and the users associated with those and quickly identifies anomalies.

For example, if botnet malware is inadvertently hosted on a well-known website looking for victims for a drive-by download, the platform still knows to evaluate the content of the file attempting to be delivered. The behavior and payload of the file can tell the platform if it is malware code being re-used, such as the case with script kiddies who re-use existing code. If the platform has seen that piece of code as part of a broader program then there's a high probability that it is malware because of our previous experience with the piece of code elsewhere. If, however, the file has no previous insight associated with it, it will be sent straight to WildFire for full analysis.

Handling Unknown Files

For sensitive government networks, unknown applications or files can be handled in two different ways:

- the file is sent to your own private instance of the WildFire Threat Intelligence Cloud, the WF-500 or
- the file is sent to the Palo Alto Networks Threat Intelligence Cloud to conduct the analysis.

Within the WildFire Threat Intelligence cloud or your own instance of the WF-500 is a custom-built hardware emulation environment, which executes the files for dynamic analysis.

WildFire determines if the file is benign or malicious. If it's malicious, WildFire declares it as malware and automatically generates signatures that are shared with all WildFire customers around the globe, and generates a detailed forensics report. Your WildFire report can show more than 100 characteristics of the file, including those above as noted under "Behavior". For example, characteristics include but are not limited to the following:

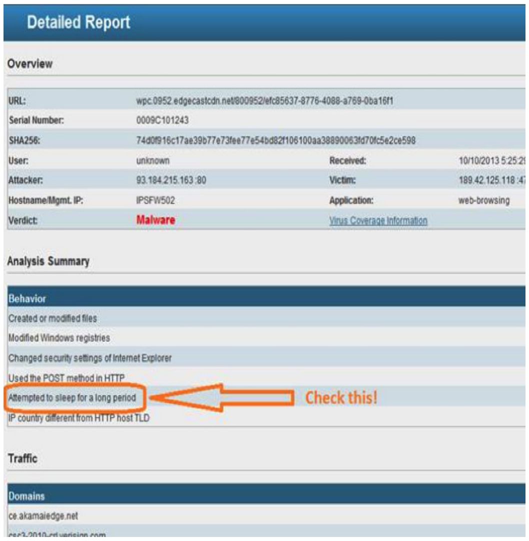

Detailed Report

Overview

URL:	wpc.0952.edgecastcdn.net/800952/efc85637-8776-4088-a769-0ba16f1		
Serial Number:	0009C101243		
SHA256:	74d0f916c17ae39b77e73fee77e54bd82f106100aa38890063fd70fc5e2ce598		
User:	unknown	Received:	10/10/2013 5:25:2
Attacker:	93.184.215.163:80	Victim:	189.42.125.118:4
Hostname/Mgmt. IP:	IPSFW502	Application:	web-browsing
Verdict:	Malware	Virus Coverage Information	

Analysis Summary

Behavior
Created or modified files
Modified Windows registries
Changed security settings of Internet Explorer
Used the POST method in HTTP
Attempted to sleep for a long period
IP country different from HTTP host TLD

Traffic

Domains
ce.akamaiedge.net
csc3-2010-crl.verisign.com

The WildFire report indicates the analysis results of the suspect file including the behaviors which led to its rating as malicious.

- It attempted to sleep for a long period
- It created and modified files
- It changed the security settings of Internet Explorer

Closed Loop Prevention

WildFire Next-generation Threat Intelligence Cloud, or the WF-500, intelligence informs the protections of Palo Alto Networks other security services, including URL filtering and Threat Prevention, for all customers. Once identified, your Palo Alto Networks security platform will detect and prevent the threat, even if your organization has never seen it before. This gives you full closed-loop threat prevention in one singular platform approach.

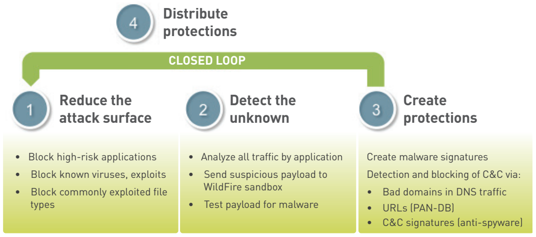

4 Distribute protections

CLOSED LOOP

1 Reduce the attack surface

- Block high-risk applications
- Block known viruses, exploits
- Block commonly exploited file types

2 Detect the unknown

- Analyze all traffic by application
- Send suspicious payload to WildFire sandbox
- Test payload for malware

3 Create protections

Create malware signatures

Detection and blocking of C&C via:

- Bad domains in DNS traffic
- URLs (PAN-DB)
- C&C signatures (anti-spyware)

Prevention Redefined: Palo Alto Networks Threat Prevention and Unknown Application and File Analysis

Prevention is still alive, thanks to Palo Alto Networks unique approach. Having detection and prevention on your network—not relying on remediation—positions your network to deal with advanced threats.

The Palo Alto Networks recognition of an unknown file type or application, transfer of the content to the WildFire Next-generation Threat Intelligence Cloud, or WF-500 for analysis, and signature creation all happens in less than 30 minutes. All the while, the Palo Alto Networks platform prevents the threat from traversing your network, leaving your network protected.

paloalto networks®
the enterprise security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

Exhibit 18

paloalto
NETWORKS

PAN-OS XML-based REST API

Usage Guide

PAN-OS 5.0

Revision B

Contents

1 Using the XML REST API

In addition to the WebUI and a Command Line Interface, PAN-OS provides a RESTful XML API to manage both the Firewall and Panorama devices. The API allows access to several types of data on the device so they can be easily integrated with and used in other systems. The API is provided as a web service that is implemented using HTTP requests and responses.

The structure of the URI for the API requests is shown below:

Beginning PAN-OS 4.1.0	http(s)://*hostname*/api/?*request-paramaters-values*
Pre-PAN-OS 4.1.0 (backward compatible in 4.1.0)	http(s)://*hostname*/esp/restapi.esp?*request-paramaters-values*

The *hostname* is the device's IP address or Domain name. The *request-parameters-values* is a series of multiple 'parameter=value' pairs separated by the ampersand character (&). The keywords for all the *parameters* are described in this document. The *values* can either be keywords or data-values in standard or XML format. The response data is always in XML format. When using the API with a command line tool such as cURL or wget, both HTTP GET and POST methods are supported.

2 API Request Types

There are currently five different API requests that can be done. These are accessed via the *type* parameter.

- Key Generation: type=*keygen*
- Device Configuration: type=*config*
- Operational Commands (PAN-OS 4.1.0 and later only): type=*op*
- Commit Configuration (PAN-OS 4.1.0 and later only): type=*commit*
- Reporting: type=*report*
- Exporting files(PCAP files supported in PAN-OS 4.1.0 and later, Other files are supported in PAN-OS 5.0.0 and later only): type=*export*
- Importingfiles (PAN-OS 5.0.0 and later only): type=*import*
- Retrieving logs (PAN-OS 5.0.0 and later only): type=*log*
- Set or Get User-ID mapping (PAN-OS 5.0.0 and later only) type=*user-id*

2.1 Key Generation

Prior to using the API, you must generate an API key that will be used for authentication for all API calls. This is done by constructing a request using credentials for an existing admin account. The API is available to all administrators (including role based) from PAN-OS 5.0.0; to only Superuser and Superuser (readonly) administrators in PAN-OS 4.1.0; and to only Superuser admins on versions before PAN-OS 4.0.0 and before. Use the URL below, replacing hostname, username, and password with the appropriate values. Any special characters in the password must be URL/percent-encoded.
http(s)://*hostname*/api/?type=keygen&user=*username*&password=*password*

The result with be an XML block that contains the key. It should look like the following:

```
<response status="success">
    <result>
        <key>gJlQWE56987nBxIqyfa62sZeRtYuIo2BgzEA9UOnlZBhU</key>
    </result>
</response>
```

The key must be URL encoded when used in HTTP requests. The API returns

From PAN-OS 4.1.0, the API returns separate keys each time a keygen query is run. All of the returned keys are valid.

To revoke or change the key, simply change the password with the associated admin account. It is recommended that you setup a new admin account to use with the API.

2.2 Device Configuration

The API allows you to configure or retrieve either all or part of the running or candidate device configuration. The API supports five options that are accessed via the *action* parameter.

- Retrieve running configuration: action=*show*
- Get candidate configuration: action=*get*
- Set candidate configuration: action=*set*
- Edit candidate configuration: action=*edit*
- Delete candidate configuration: action=*delete*
- Rename a configuration object: action=*rename*
- Clone a configuration object: action=*clone*
- Move a configuration object: action=*move*

2.2.1 Retrieve

Using *action=show* with no additional parameters, will return the entire running configuration. Using the *xpath* parameter, you target a specific portion of the configuration. For example, to retrieve just the security rulebase, use: xpath=/config/devices/entry/vsys/entry/rulebase/security. **NOTE**: There is no trailing backslash character at the end of the *xpath*.

The URL for the API request will be:
http(s)://*hostname*/api/?type=config&action=show&key=*keyvalue*&xpath=/config/devices/entry/vsys/entry/rulebase/security

The XML response for the query should look like the following (truncated):

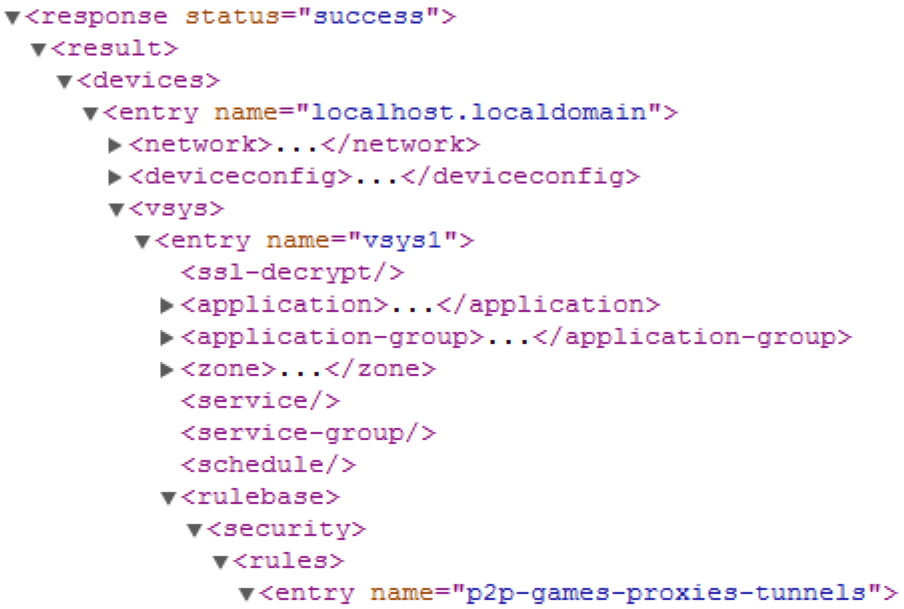

```
▼<response status="success">
  ▼<result>
    ▼<devices>
      ▼<entry name="localhost.localdomain">
        ▶<network>...</network>
        ▶<deviceconfig>...</deviceconfig>
        ▼<vsys>
          ▼<entry name="vsys1">
              <ssl-decrypt/>
            ▶<application>...</application>
            ▶<application-group>...</application-group>
            ▶<zone>...</zone>
              <service/>
              <service-group/>
              <schedule/>
            ▼<rulebase>
              ▼<security>
                ▼<rules>
                  ▼<entry name="p2p-games-proxies-tunnels">
```

2.2.2 Get

Beginning with PAN-OS 4.1.0, you can get the candidate configuration from the firewall or Panorama device using the Config Get API request. Use the *xpath* parameter to specify the portion of the configuration to get.

http(s)://hostname/api/?type=config&action=get&xpath=*path-to-config-node*

For instance to get the address objects in a VSYS, you can use the following:
http(s)://hostname//api/?type=config&action=get&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address

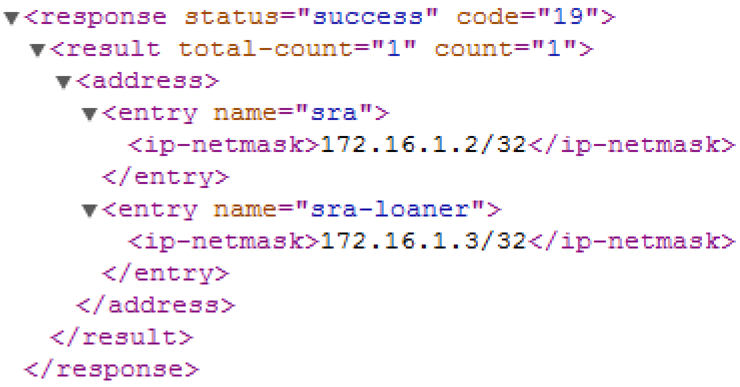

```
▼<response status="success" code="19">
  ▼<result total-count="1" count="1">
    ▼<address>
      ▼<entry name="sra">
          <ip-netmask>172.16.1.2/32</ip-netmask>
        </entry>
      ▼<entry name="sra-loaner">
          <ip-netmask>172.16.1.3/32</ip-netmask>
        </entry>
      </address>
    </result>
  </response>
```

To get the pre-rules pushed from panorama, you can use the following:
http(s)://firewall//api/?type=config&action=get&xpath=/config/panorama/vsys/entry[@name='vsys']/pre-rulebase/security

You can use this query is to get detail information on Applications and Threats from the firewall.
http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/threats/vulnerability/entry[@name='30003']

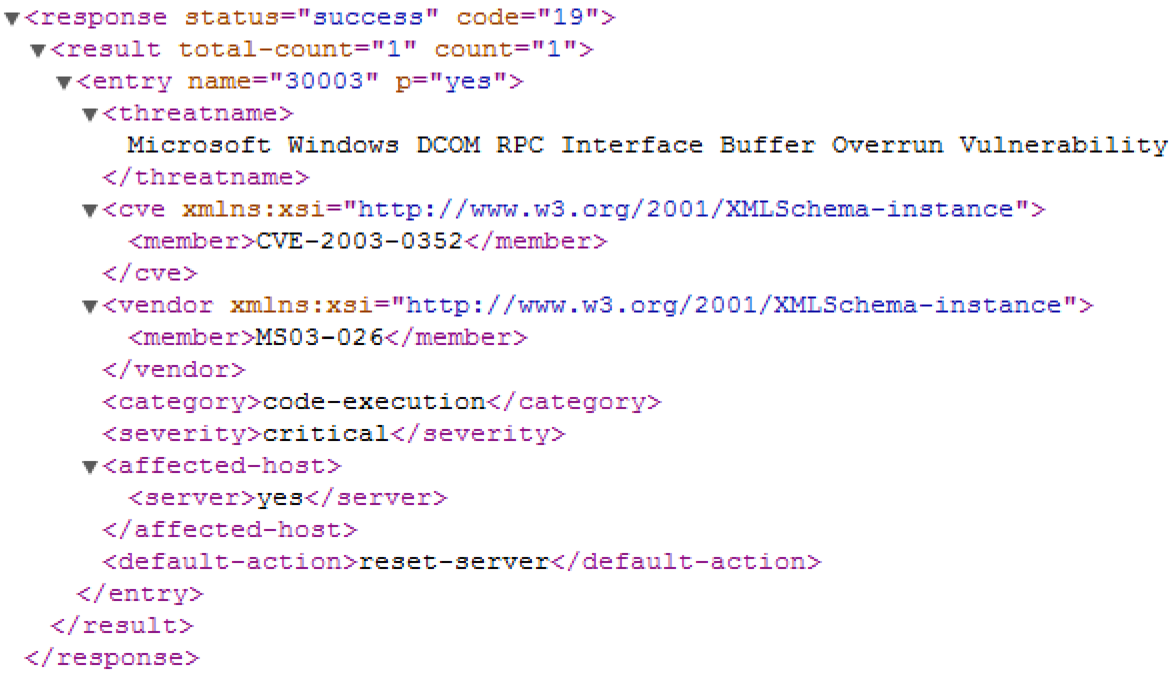

```
▼<response status="success" code="19">
  ▼<result total-count="1" count="1">
    ▼<entry name="30003" p="yes">
      ▼<threatname>
          Microsoft Windows DCOM RPC Interface Buffer Overrun Vulnerability
        </threatname>
      ▼<cve xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
          <member>CVE-2003-0352</member>
        </cve>
      ▼<vendor xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
          <member>MS03-026</member>
        </vendor>
        <category>code-execution</category>
        <severity>critical</severity>
      ▼<affected-host>
          <server>yes</server>
        </affected-host>
        <default-action>reset-server</default-action>
      </entry>
    </result>
  </response>
```

http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/application, provides details on the full list of all applications.
http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/application/entry[@name='hotmail'], provides details on the specific application.

Refer to the API browser and follow the 'Configuration Commands' link to see all the available config xpaths.

2.2.3 Set

Using *action=set*, you can add or create a new object at a specified location in the configuration hierarchy. Use the *xpath* parameter to specify the location of the object in the configuration and the *element* parameter to specify a value for the object using its XML representation (as seen in the output of action=show).

For instance, to create a new rule called *rule1* in the security policy, use the below Config Set API request:
http(s)://hostname/api/?type=config&action=set&key=keyvalue&xpath=*xpath-value*&element=*element-value,* where

xpath-value is /config/devices/entry/vsys/entry/rulebase/security/rules/entry[@name='rule1'], and
element-value is
<source><member>src</member></source><destination><member>dst</member></destination><service><member>service</member></service><application><member>application</member></application><action>action</action><source-user><member>src-user</member></source-user><option><disable-server-response-inspection>yes-or-no</disable-server-response-inspection></option><negate-source>yes-or-no</negate-source><negate-destination>yes-or-no</negate-destination><disabled>yes-or-no</disabled><log-start>yes-or-no</log-start><log-end>yes-or-no</log-end><description>description</description><from><member>src-zone</member></from><to><member>dst-zone</member></to>

Use the response from the config show API request to create the xml body for the element.
http(s)://hostname/api/?type=config&action=show

To add an additional member to a group, use member[text()='name'] in the xpath. For e.g., to add an additional address object named *abc* to a address group named *test*, use:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address-group/entry[@name='test']&element=<member>abc</member>

2.2.4 Edit

Using action=*edit*, you can replace an existing object hierarchy at a specified location in the configuration with a new value. Use the *xpath* parameter to specify the location of the object and the *element* parameter to specify a new value for the object using its XML object hierarchy (as seen in the output of action=show). For instance, to replace the application(s) currently used in a rule *rule1* with a new application, use:
http(s)://hostname/api/?type=config&action=edit&key=keyvalue&xpath=*xpath-value*&element=*element-value,* where

xpath=/config/devices/entry/vsys/entry/rulebase/security/rules/entry[@name='rule1']/application
element=<application><member>app-name</member></application>

Use the response from the config show API request to create the xml body for the element.
http(s)://hostname/api/?type=config&action=show

To replace all members in a node with a new set of members, use the entry tag in both the xpath and element parameters. For e.g., to replace all the address objects in the address group named test with two new members named *abc* and *xyz*, use:
http(s)://hostname/api/?type=config&action=edit&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address-group/entry[@name='test']&element=<entry name='test'><member>abc</member><member>xyz</member></entry>

2.2.5 Delete

Using action=*delete*, you can delete an object at a specified location in the configuration. Use xpath parameter to specify the location of the object to be deleted. For instance, to delete a rule named *rule1* in the security policy, use the below API query:
http(s)://hostname/api/?type=config&action=delete&xpath=/config/devices/entry/vsys/entry/rulebase/security/rules/entry[@name='rule1']
To delete a single member object in a group, use the object name in the xpath as member[text()='name']. For e.g., to delete an address object named *abc* in an address group named *test*, use the below xpath:
http(s)://hostname/api/?type=config&action=delete&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address-group/entry[@name='test']/member[text()='abc']

2.2.6 Rename

Using action=*rename*, you can rename an object at a specified location in the configuration. Use the *xpath* parameter to specify the location of the object to be renamed. Use the *newname* parameter to provide a new name for the object.
For instance, to rename an address object called old_address to new_address, use the below API query:
http(s)://hostname/api/?type=config&action=*rename*&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address/entry[@name='old_address']&*newname*=new_address

2.2.7 Clone

Using action=*clone*, you can clone an existing configuration object. Use the *xpath* parameter to specify the location of the object to be cloned. Use the '*from*' parameter to specify the source object, and the *newname* parameter to provide a name for the cloned object. For instance, to clone a security policy called rule1 into rule2, use the below API query:
http(s)://hostname/api/?type=config&action=*clone*&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/rulebase/security/rules/&*from*=/config/devices/entry/vsys/entry[@name='vsys1']/rulebase/security/rules/entry[@name='rule1']&*newname*=rule2

2.2.8 Move

Using action=*move*, you can move the location of an existing configuration object. Use the *xpath* parameter to specify the location of the object to be moved, the *where* parameter to specify type of move, and *dst* parameter to specify the destination xpath.

- where=*after*&dst=xpath
- where=*before*&dst=xpath
- where=top
- where=bottom

For instance, to move a security policy called rule1 after rule2, use the below API query:
http(s)://hostname/api/?type=config&action=*move*&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/rulebase/security/rules/entry[@name='rule1']&where=after&dst=rule2

2.3 Commit

Beginning with PAN-OS 4.1.0, you can commit candidate configuration to a firewall or Panorama device using the commit API request.

To commit a candidate configuration, use the following:
http(s)://hostname/api/?type=commit&cmd=<commit></commit>

To do a force commit of the candidate configuration, use the following:
http(s)://hostname/api/?type=commit&cmd=<commit><force>*body*</force></commit>

To do a granular or partial commit of the candidate configuration, use the following:

http(s)://hostname/api/?type=commit&cmd=<commit><partial>*body*</partial></commit>

Refer to the API browser for the different options available for use with force and partial commits. The *body* element in the *cmd* parameter should be replaced by the XML element for the corresponding commit operation.

When there are no pending changes to commit, API request returns:

```
<response status="success" code="19">
        <msg>There are no changes to commit.</msg>
</response>
```

When there are pending changes, the API returns a Job ID for the commit request as below.

```
<response status="success" code="19">
        <result>
                <msg><line>Commit job enqueued with jobid 4</line></msg>
                <job>4</job>
        </result>
</response>
```

You can query the status of the job using the below Operational API request and the corresponding response:

```
http(s)://hostname/api/?type=op&cmd=<show><jobs><id>4</id></jobs></show>
<response status="success">
        <result>
                <job>
                        <tenq>2011/10/20 20:41:44</tenq>
                        <id>4</id>
                        <type>Commit</type>
                        <status>FIN</status>
                        <stoppable>no</stoppable>
                        <result>OK</result>
                        <tfin>20:42:22</tfin>
                        <progress>20:42:22</progress>
                        <details><line>Configuration committed successfully</line></details>
                        <warnings/>
                </job>
        </result>
</response>
```

2.3.1 Commit-all (Panorama)

Beginning with PAN-OS 4.1.0, you can push shared policy from Panorama to all centrally managed firewall devices using the commit-all API request type

To push configuration to an entire device group (say *west-dg*), use the following:
http(s)://panorama/api/?type=commit&action=all&cmd=<commit-all><shared-policy><device-group>*west-dg*</device-group></shared-policy></commit-all>

To push configuration to a VSYS (say *mktg-vsys*), use the following:
http(s)://panorama/api/?type=commit&action=all&cmd=<commit-all><shared-policy><vsys>***mktg-vsys***</vsys></shared-policy></commit-all>

To push configuration to a specific device by serial number, use the following:

http(s)://panorama/api/?type=commit&action=all&cmd=<commit-all><shared-policy><device></device></shared-policy></commit-all>

Refer to the API browser for other options available for granular commit operations on Panorama. The values for the *cmd* parameter should be replaced by the XML element for the corresponding commit operation.

2.4 Operational Commands

Beginning with PAN-OS 4.1.0, you can use any of the operational commands available on the command line interface using the *Op* API request below:
http(s)://hostname/api/?type=op&cmd=*xml-body*

Refer to the API browser and follow the link for operational commands to see a complete listing of all the different options available for the *xml-body* and their corresponding operation.

Examples of operational API requests include setting, showing, or clearing runtime parameters, saving and loading configurations to disk, retrieving interface or system information, etc.

To request a system restart, use:
http(s)://hostname/api/?type=op&cmd=<request><restart><system></system></restart></request>

To install system software version 4.1.0, use:
http(s)://hostname/api/?type=op&cmd=<request><system><software><install><version>4.1.0</version></install></software></system></request>

To set the system setting to turn on multi-vsys mode, use:
http(s)://hostname/api/?type=op&cmd=<set><system><setting><multi-vsys></multi-vsys></setting></system></set>

To schedule a User Activity Report, use:
http(s)://hostname/api/?type=op&cmd=<schedule><uar-report><user>*username*</user><title>*titlename*</title></uar-report></schedule>

To save or load config to/from a file, use:
http(s)://hostname/api/?type=op&cmd=<save><config><to>*filename*</to></config></save>, and
http(s)://hostname/api/?type=op&cmd=<load><config><from>*filename*</from></config></load>

2.5 Reporting

The XML API provides a way to quickly pull the results of any report defined in the system using the type=*report* parameter. There are three report stores that can be pulled from:

- Dynamic Reports (ACC reports): reporttype=*dynamic*
- Predefined Reports: reporttype=*predefined*
- Custom Reports: reporttype=*custom*

To retrieve a specific report by name, use the *reportname* parameter:
http(s)://hostname/api/?type=report&reporttype=*dynamic*|*predefined*|*custom*&reportname=*name*

Note: When generating an on-demand report on Panorama, when you use the query, the on-screen output will display a job-id instead of the requested report. To retrieve the report, you must poll the job status using the job id (shown in 2.6.3) until the job completes. On completion, the job status reports as FIN (finished), and the reports displays.

2.5.1 Dynamic reports

The names for the currently supported *dynamic* reports follows:

- acc-summary
- custom-dynamic-report
- top-app-summary
- top-application-categories-summary
- top-application-risk-summary
- top-application-subcategories-summary
- top-application-tech-summary
- top-applications-summary
- top-applications-trsum
- top-attacker-countries-summary
- top-attackers-summary
- top-attacks-acc
- top-blocked-url-categories-summary
- top-blocked-url-summary
- top-blocked-url-user-behavior-summary
- top-data-dst-countries-summary
- top-data-dst-summary
- top-data-egress-zones-summary
- top-data-filename-summary
- top-data-filetype-summary
- top-data-ingress-zones-summary
- top-data-src-countries-summary
- top-data-src-summary
- top-data-type-summary
- top-dst-countries-summary
- top-dst-summary
- top-egress-zones-summary
- top-hip-objects-details
- top-hip-objects-summary
- top-hip-profiles-details
- top-hip-profiles-summary
- top-hip-report-links
- top-hr-applications-summary
- top-ingress-zones-summary
- top-rule-summary
- top-spyware-phonehome-summary
- top-spyware-threats-summary
- top-src-countries-summary
- top-src-summary
- top-threat-egress-zones-summary
- top-threat-ingress-zones-summary
- top-threats-type-summary
- top-url-categories-summary
- top-url-summary
- top-url-user-behavior-summary
- top-victim-countries-summary
- top-victims-summary
- top-viruses-summary
- top-vulnerabilities-summary

You can get the above list of dynamic report names using the below API request, or by following the links on the API browser. http(s)://hostname/api/?type=*report*&reporttype=*dynamic*

For dynamic reports, you can provide the timeframe for the report via the *period* or *starttime* and *endtime* options(use a + instead of a space between the date and timestamp). The number of rows is set via *topn*. The possible values for *period* are:

- last-60-seconds
- last-15-minutes
- last-hour
- last-12-hrs
- last-24-hrs
- last-calendar-day
- last-7-days
- last-7-calendar-days
- last-calendar-week
- last-30-days

2.5.2 Predefined reports

The names for the currently supported *predefined* reports are shown below. Predefined reports always return data for the last 24 hour period. You can also get this list by following the link for predefined reports on the API browser or running this API query: http(s)://hostname/api/?type=*report*&reporttype=*predefined*

- bandwidth-trend
- risk-trend
- risky-users
- spyware-infected-hosts
- threat-trend
- top-application-categories
- top-applications
- top-attackers
- top-attackers-by-countries
- top-attacks
- top-blocked-url-categories
- top-blocked-url-user-behavior
- top-blocked-url-users
- top-blocked-websites
- top-connections
- top-denied-applications
- top-denied-destinations
- top-denied-sources
- top-destination-countries
- top-destinations
- top-egress-interfaces
- top-egress-zones
- top-http-applications
- top-ingress-interfaces
- top-ingress-zones
- top-rules
- top-source-countries
- top-sources
- top-spyware-threats
- top-technology-categories
- top-url-categories
- top-url-user-behavior
- top-url-users
- top-users
- top-victims
- top-victims-by-countries
- top-viruses
- top-vulnerabilities
- top-websites
- unknown-tcp-connections
- unknown-udp-connections

2.5.3 Custom reports

For custom reports, the different selection criteria (time frame, group-by, sort-by, etc.) are part of the report definition itself. The API returns any shared custom reports. Note that quotes are not required around the reportname and any spaces in the report name must be URL encoded to %20.

For custom reports created in a specific VSYS, you can retrieve them directly by specifying the vsys parameters. This functionality is only available beginning PAN-OS 4.1.1. Prior to PAN-OS 4.1.1., you will need to follow a 2-step process.

Step one, retrieve the report definition from the configuration using a Config Get API request. For e.g., a report named *report-abc*:
http(s)://firewall/api/?type=config&action=get&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/reports/entry[@name='*report-abc*']

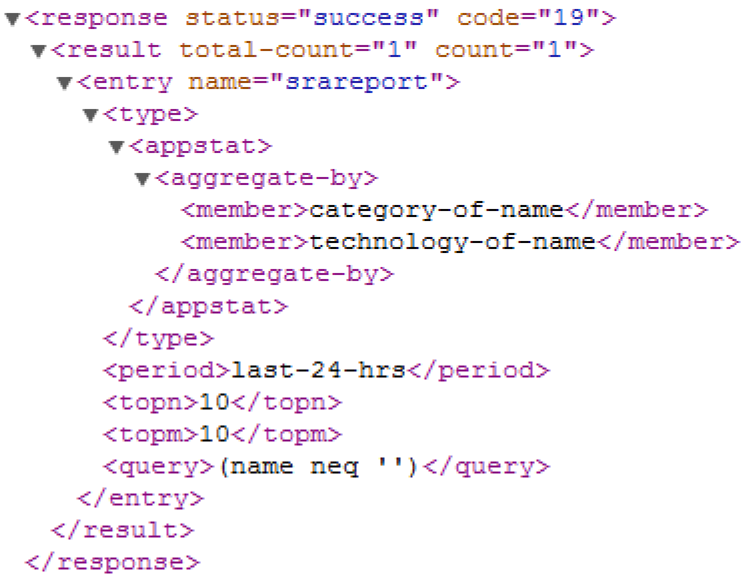

```
▼<response status="success" code="19">
  ▼<result total-count="1" count="1">
    ▼<entry name="srareport">
      ▼<type>
        ▼<appstat>
          ▼<aggregate-by>
              <member>category-of-name</member>
              <member>technology-of-name</member>
            </aggregate-by>
          </appstat>
        </type>
        <period>last-24-hrs</period>
        <topn>10</topn>
        <topm>10</topm>
        <query>(name neq '')</query>
      </entry>
    </result>
  </response>
```

Step Two, retrieve a dynamic report data using reporttype=*dynamic*, reportname=*custom-dynamic-report*, and cmd=*report-definition* where report definition is the XML definition retrieved in the previous query.
http(s)://hostname/api/?type=report&reporttype=dynamic&reportname=custom-dynamic-report&cmd=<type><appstat><aggregate-by><member>category-of-name</member><member>technology-of-name</member></aggregate-by></appstat></type><period>last-24-hrs</period><topn>10</topn><topm>10</topm><query>(name neq '') AND (vsys eq ‘vsys1’)</query>

2.6 Exporting files

You can export certain types of files from the firewall using the type=*export* parameter in the API request. The type of file to be exported must be specified using the *category* parameter.
From PAN-OS 4.1.0 onwards for:

- Packet Captures: category=<*application-pcap* | *threat-pcap* | *filter-pcap* | *filters-pcap*>

From PAN-OS 5.0.0 onwards for:

- Configuration: category=*configuration*
- Certificates/Keys: category=<*certificate* | *high-availability-key* | *key-pair*>
- Response pages: category=< *application-block-page* | *captive-portal-text* | *file-block-continue-page* | *file-block-page* | *global-protect-portal-custom-help-page* | *global-protect-portal-custom-login-page* | *global-protect-portal-custom-welcome-page* | *ssl-cert-status-page* | *ssl-optout-text* | *url-block-page* | *url-coach-text* | *virus-block-page*>
- Technical support data: category=*tech-support*
- Device State: category=*device-state*

Use wget or cURL tools to export the file from the firewall and save locally with a local file name, as below. Refer to their respective man pages for additional usage information.

>wget --output-document=*filename* “http(s)://firewall/api/?*query-parameters*”

>curl -o *filename* “http(s)://firewall/api/?*query-parameters*”

When using the API query from a web-browser, you can specify to=filename as an optional parameter if you would like to provide a different name when saving the file locally.

2.6.1 Packet Captures

You can export packet captures from the firewall device using the Export API request. The type of PCAP to be exported using the API must be specified using the *category* parameter.

- Application Packet Captures: category=*application-pcap*
- Threat Packet Captures: category=*threat-pcap*
- Debug Filter Packet Captures: category=*filter-pcap* or *filters-pcap*
- Data filtering Packet Captures: category=*dlp-pcap*. This requires a dlp-password parameter to work.

2.6.1.1 Application and Threat PCAPs

Application and Threat PCAPs are organized by a *Directory/Filename* structure where the directory is a date in yyyymmdd format. Filename for application pcaps uses a SourceIP-SourcePort-DestinationIP-DestinationPort-SessionID.pcap format. Filename for threat pcaps uses a Epoch-SessionID.pcap format.

To get a list of directories for the application and threat PCAPs, you can use the following:
http(s)://firewall/api/?type=export&category=application-pcap, and
http(s)://firewall/api/?type=export&category=threat-pcap

To get a list of file names under a directory for the application and threat PCAPs, you can use the *from* parameter as follows:
http(s)://firewall/api/?type=export&category=application-pcap&from=*yyyymmdd*, and
http(s)://firewall/api/?type=export&category=threat-pcap&from=*yyyymmdd*

To retrieve a specific application or threat PCAP file by its name, you can use the *from* parameter as below:
http(s)://firewall/api/?type=export&category=application-pcap&from=*yyyymmdd/filename*, and
http(s)://firewall/api/?type=export&category=threat-pcap&from=*yyyymmdd/filename*
The file will be retrieved and saved locally using the name *yyyymmdd-filename*.

To retrieve a specific application or threat PCAP file by its name, and save it locally by a custom name, you can use the *to* parameter as below:
http(s)://firewall/api/?type=export&category=application-pcap&from=*yyyymmdd/filename*&to=*localfile*, and
http(s)://firewall/api/?type=export&category=threat-pcap&from=*yyyymmdd/filename*&to=*localfile*

2.6.1.2 Debug filter PCAPs

To get a list of filter PCAP file names, you can use:
http(s)://firewall/api/?type=export&category=filters-pcap

To retrieve a specific filter PCAP file, you can use:
http(s)://firewall/api/?type=export&category=filters-pcap&from=*filename*

2.6.1.3 Data filtering PCAPs

To get a list of data filtering PCAP file names, you can use:
http(s)://firewall/api/?type=export&category=dlp-pcap&dlp-password=<password>

To retrieve a specific data filtering PCAP file, you can use:
http(s)://firewall/api/?type=export&category=dlp-pcap&dlp-password=<password>&from=*filename*&to=<localfile>

2.6.2 Certificates/Keys

There are additional query parameters to be specified when exporting Certificates/Keys from the firewall.
http(s)://firewall/api/?type=export&category=*certificate*&certificate-name=<certificate_name>&format=<*pkcs12* | *pem*>&include-key=<*yes* | *no*>&vsys=<vsys | omit this parameter to import it into shared location>

- certificate-name: name of the certificate object on the firewall

- format: cerficate format, pkcs12 or pem
- include-key: yes or no parameter to include or exclude the key
- passphrase: required when including the certificate key
- vsys: Virtual System where the certificate object is used. Ignore this parameter if the certificate is a shared object.

2.6.3 Technical Support Data (debug logs etc.)

Since debug log data sizes are large, the API uses an asynchronous job scheduling approach to retrieve technical support data. The initial query creates a Job id with a hash that is used in the follow on queries with the action parameter. The values for the action parameter are:

- When action parameter is not specified, the system creates a new job to retrieve tech support data.
- action=*status*, to check status of the job. Returns either PEND or FIN.
- action=*get*, to retrieve the data when the status shows FIN.
- action=*finish*, to manually delete the job.

Create a job to retrieve technical support data using http(s)://firewall/api/?type=export&category=tech-support, which returns a job id.

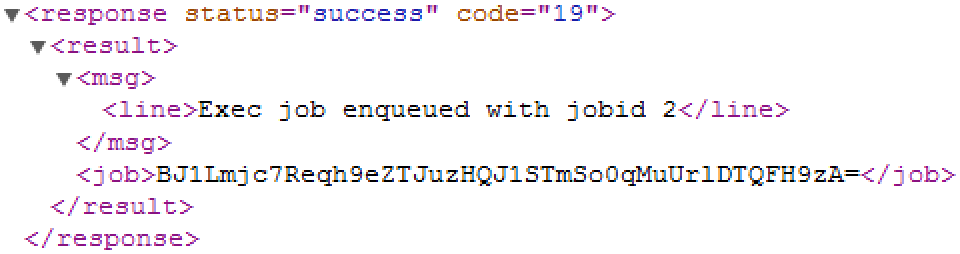

```
▼<response status="success" code="19">
  ▼<result>
    ▼<msg>
       <line>Exec job enqueued with jobid 2</line>
     </msg>
     <job>BJ1Lmjc7Reqh9eZTJuzHQJ1STmSo0qMuUr1DTQFH9zA=</job>
   </result>
 </response>
```

Check the status of the job using: http(s)://firewall/api/?type=*export*&category=*tech-support*&action=*get*&job-id=*id*. Use the job id returned in the previous response as the job-id parameter. A status value of ‘FIN’ indicates the data is ready to be retrieved.

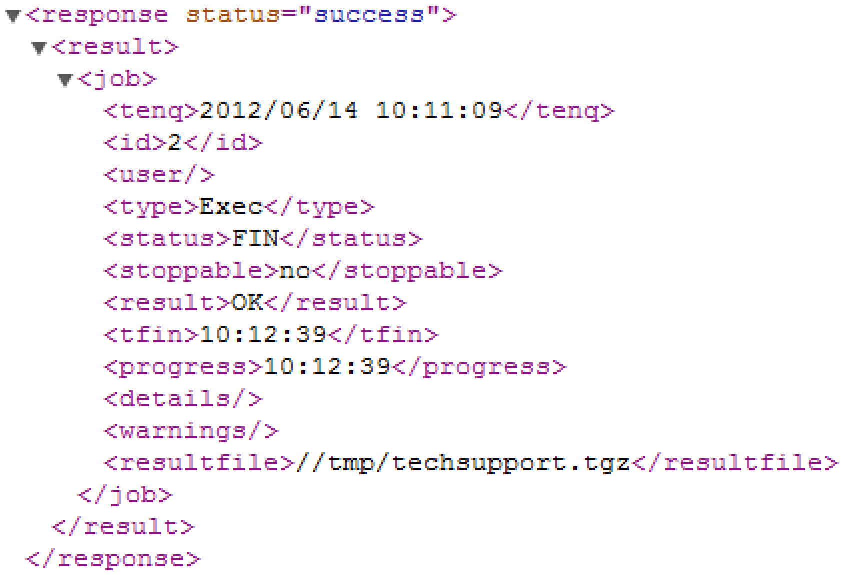

```
▼<response status="success">
  ▼<result>
    ▼<job>
       <tenq>2012/06/14 10:11:09</tenq>
       <id>2</id>
       <user/>
       <type>Exec</type>
       <status>FIN</status>
       <stoppable>no</stoppable>
       <result>OK</result>
       <tfin>10:12:39</tfin>
       <progress>10:12:39</progress>
       <details/>
       <warnings/>
       <resultfile>//tmp/techsupport.tgz</resultfile>
     </job>
   </result>
 </response>
```

Retrieve the data using: http(s)://firewall/api/?type=*export*&category=*tech-support*&action=*get*&job-id=*id*. When using cURL or wget, you can specify the output file name as an option to cURL (-o) or wget (--output-document). After a successful retrieval of the job data, the job is automatically deleted by the system.

To manually delete the job use: http(s)://firewall/api/?type=*export*&category=*tech-support*&action=*finish*&job-id=*id*

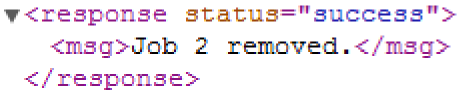

```
▼<response status="success">
   <msg>Job 2 removed.</msg>
 </response>
```

2.7 Importing files

Beginning with PAN-OS 5.0.0, you can import certain types of files into the firewall using the type=*import* parameter in the API request. The type of file to be imported must be specified using the *category* parameter.

- Software: category=*software*
- Content: category=<*anti-virus* | *content* | *url-database* | *signed-url-database*>
- Licenses: category=*license*
- Configuration, category=*configuration*
- Certificates/Keys, category=<*certificate* | *high-availability-key* | *key-pair*>
- Response pages, category=< *application-block-page* | *captive-portal-text* | *file-block-continue-page* | *file-block-page* | *global-protect-portal-custom-help-page* | *global-protect-portal-custom-login-page* | *global-protect-portal-custom-welcome-page* | *ssl-cert-status-page* | *ssl-optout-text* | *url-block-page* | *url-coach-text* | *virus-block-page*>
- Clients, category=*global-protect-client*
- Custom logo, category=*custom-logo*

Use wget or cURL tools to import the file to the firewall, as below. Refer to their respective man pages for additional usage information.

>wget --post-file *filename* "http(s)://firewall/api/?*query-parameters*&client=*wget* &file-name=*filename*"

>curl --form file=@*filename* "http(s)://firewall/api/?*query-parameters*"

2.7.1 Certificates/Keys

There are additional query parameters to be specified when importing Certificates/Keys to the firewall. The type of the certificate or key file is specified using the category parameter

- category=*certificate*
- category=*keypair*
- category=*high-availability-key*

The certificate file import (category=*certificate*) and keypair import (category=*keypair*) take the below additional parameters.

- certificate-name: name of the certificate object on the firewall
- format: cerficate format, pkcs12 or pem
- passphrase: required when including the certificate key
- vsys: Virtual System where the certificate object is used. Ignore this parameter if the certificate is a shared object.

For e.g., http(s)://firewall/api/?type=*import*&category=*certificate*&certificate-name=<certificate_name>&format=<*pkcs12* | *pem*>&passphrase=<text>&vsys=<vsys | omit this parameter to import it into shared location>

2.7.2 Response pages

Only the GlobalProtect related response pages require an additional parameter for the *profile* where the page should be imported to.

- profile=*profilename*

2.7.3 Custom logo

Custom logos can be imported to different locations based on the where parameter.

- where=<*login-screen* | *main-ui* | *pdf-report-footer* | *pdf-report-header*>

2.8 Retrieving Logs

Beginning with PAN-OS 5.0.0, you can retrieve logs from the firewall using the API with the type=*log* parameter. The type of logs to retrieve must be specified using the log-type parameter.

- log-type=*traffic*, for traffic logs
- log-type=*threat*, for threat logs,
- log-type=*config*, for config logs,
- log-type=*system*, for system logs,
- log-type=*hip-match*, for HIP logs.

The other optional parameters to this request are:

- *query* parameter to specify match criteria for the logs. This is similar to the query provided in the WebUI under the Monitor tab when viewing the logs. The query must be URL encoded.
- *nlogs* parameter to specify the number of logs to be retrieved. The default is 20 when the parameter is not specified. The maximum is 5000.
- *skip* parameter to specify the number of logs to skip when doing a log retrieval. The default is 0. This is useful when retrieving logs in batches where you can skip the previously retrieved logs.

Since log data sizes can be large, the API uses an asynchronous job scheduling approach to retrieve log data. The initial query returns a Job id with a Hash that is used in the follow on queries with the action parameter. The values for the action parameter are:

- Unspecified: when the action parameter is not specified, the system creates a new job to retrieve log data.
- action=*get*, to check status and retrieve the log data when the status is FIN. (This is a slight difference from the asynchronous approach to retrieve tech support data where a separation status action was available)
- action=*finish*, to manually delete the job.

To create a job to retrieve all traffic logs that occurred after a certain time, you can use below query. NOTE: A web-browser will automatically URL encode the parameters, but when using wget/curl tools, the query parameter must be URL encoded.

http(s)://firewall/api/?type=log&log-type=traffic&query=(receive_time geq '2012/06/22 08:00:00')

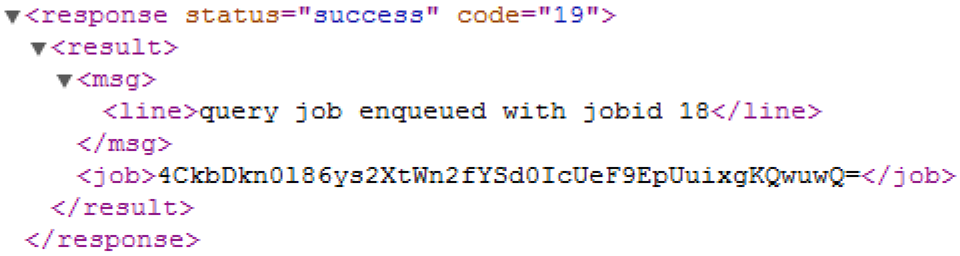

```
▼<response status="success" code="19">
  ▼<result>
    ▼<msg>
       <line>query job enqueued with jobid 18</line>
     </msg>
     <job>4CkbDkn0186ys2XtWn2fYSd0IcUeF9EpUuixgKQwuwQ=</job>
   </result>
 </response>
```

Retrieve the data using: http(s)://firewall/api/?type=log&action=get&job-id=*id*, where id is the value returned in the previous response.

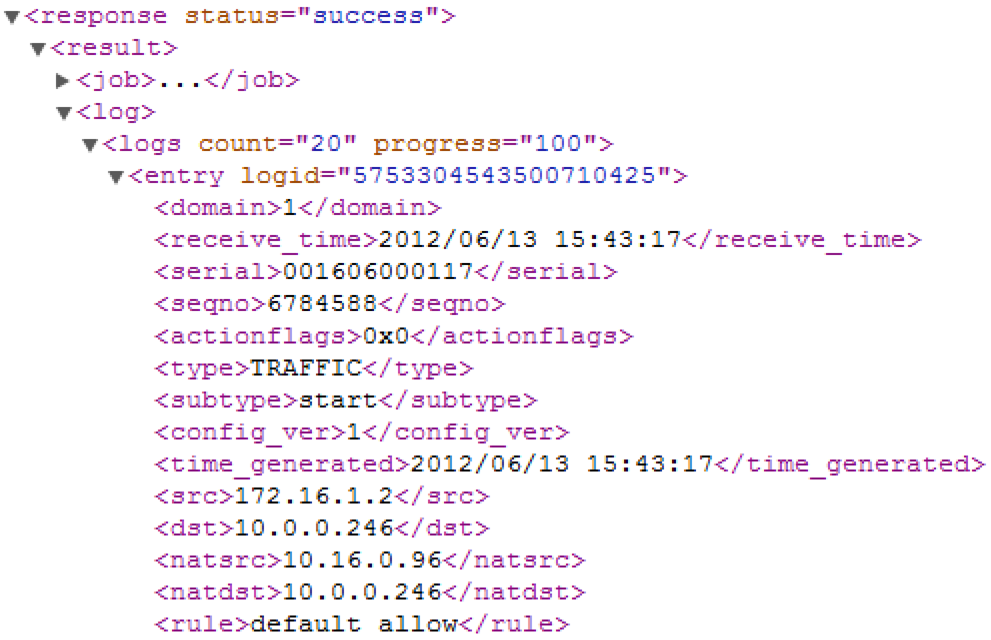

```
▼<response status="success">
  ▼<result>
    ▶<job>...</job>
    ▼<log>
      ▼<logs count="20" progress="100">
        ▼<entry logid="5753304543500710425">
            <domain>1</domain>
            <receive_time>2012/06/13 15:43:17</receive_time>
            <serial>001606000117</serial>
            <seqno>6784588</seqno>
            <actionflags>0x0</actionflags>
            <type>TRAFFIC</type>
            <subtype>start</subtype>
            <config_ver>1</config_ver>
            <time_generated>2012/06/13 15:43:17</time_generated>
            <src>172.16.1.2</src>
            <dst>10.0.0.246</dst>
            <natsrc>10.16.0.96</natsrc>
            <natdst>10.0.0.246</natdst>
            <rule>default allow</rule>
```

When the job status is FIN (finished), the response automatically includes all the logs in the xml data response. The <log> node in the xml data is not present when the job status is still pending. After successful log data retrieval, the system automatically deletes the job.

To manually delete a log retrieval job, you must run the below query.
http(s)://firewall/api/?type=log&action=finish&job-id=*id*, which on successful completion returns:

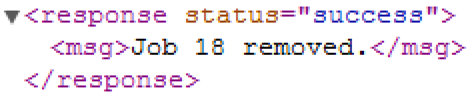

```
▼<response status="success">
    <msg>Job 18 removed.</msg>
  </response>
```

2.9 User-ID mapping

Beginning with PAN-OS 5.0.0, you can apply User-ID mapping information directly to the firewall using the API with the type=*user-id* parameter. Additionally you can also register a Dynamic Address object using this API request. It takes the following parameters.

- action=set
- Input file containing the User-ID mapping information.

>wget --post-file *filename* “http(s)://firewall/api/?type=*user-id*&action=*set*&client=*wget* &file-name=*filename*”

>curl --form file=@*filename* “http(s)://firewall/api/? type=*user-id*&action=*set*”

When providing a User-ID mapping for a login event, logout event, or for groups, the input file format is as shown below.

```
<uid-message>
    <version>1.0</version>
    <type>update</type>
    <payload>
        <login>
            <entry name="domain\uid1" ip="10.1.1.1" timeout="20">
                <hip-report>
```

```
                        .....

                </hip-report>
            </entry>
        </login>
        <groups>
            <entry name="group1">
                <members>
                    <entry name="domain\user1"/>
                    <entry name="domain\user2"/>
                </members>
            </entry>
            <entry name="group2">
                <members>
                    <entry name="domain\user3"/>
                </members>
            </entry>
        </groups>
    </payload>
</uid-message>
```

When registering an IP address for a Dynamic Address Objects, the input file format is as shown below.

```
<uid-message>
    <version>1.0</version>
    <type>update</type>
    <payload>
        <register>
            <entry identifier="CBB09C3D-3416-4734-BE90-0395B7598DE3" ip="10.1.1.1"/>
            <entry identifier="CBB09C3D-3416-4734-BE90-0395B7598DE4" ip="10.1.1.2"/>
        </register>
        <unregister>
            <entry identifier="CBB09C3D-3416-4734-BE90-0395B7598DE5" ip="10.1.1.3"/>
        </unregister>
    </payload>
</uid-message>
```

3 Panorama to device redirection

You can use the API on the Panorama to redirect the queries to a specific firewall device managed by the Panorama using the target parameter. The target parameter takes the device serial number as a value. For instance, to run a Panorama query that directs an operational command to a firewall device, you can use.
http(s)://panorama/api/?type=op&cmd=<show><system><info></info></system></show>&target=device-serial-number

4 Targeting a specific Virtual System

Use the vsys parameter to target the API request to a specific Virtual System. You can use this parameters for all Operational commands, Dynamic reports, Custom reports, and User-ID. For configuration commands, the xpath for virtual system specific objects includes the virtual system. For e.g. the xpath for an address group object in vsys1 is /config/devices/entry/vsys/entry[@name='vsys1']/address-group/entry[@name='test'].

5 Admin Access Rights

The different Administrators and Admin roles supported in the API is listed in the table below.

Version	
PAN-OS 5.0.0 and later	*Dynamic roles*: Superuser, Superuser (readonly), Device admin, Device admin (readonly), Vsys admin, Vsys admin (readonly) *Role based Admins*: Device, Vsys, Panorama.
PAN-OS 4.1.0	*Dynamic roles*: Superuser, Superuser (readonly)
PAN-OS 4.0.0 and older	*Dynamic roles*: Superuser

For Admin role profiles, permissions can be enabled or disabled on the basis of the *type* parameter as below.

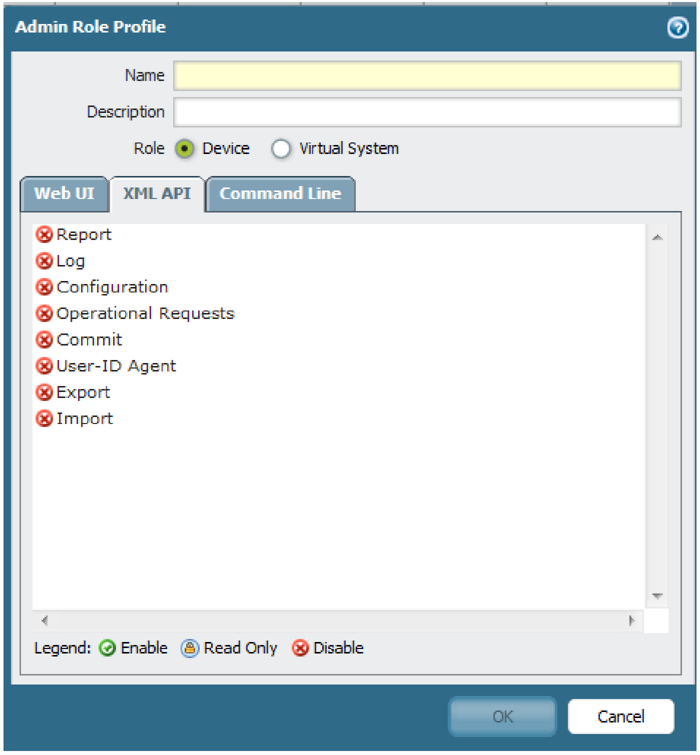

Name

Description

Role ● Device ○ Virtual System

⊗ Report
⊗ Log
⊗ Configuration
⊗ Operational Requests
⊗ Commit
⊗ User-ID Agent
⊗ Export
⊗ Import

Legend: ✓ Enable ◉ Read Only ⊗ Disable

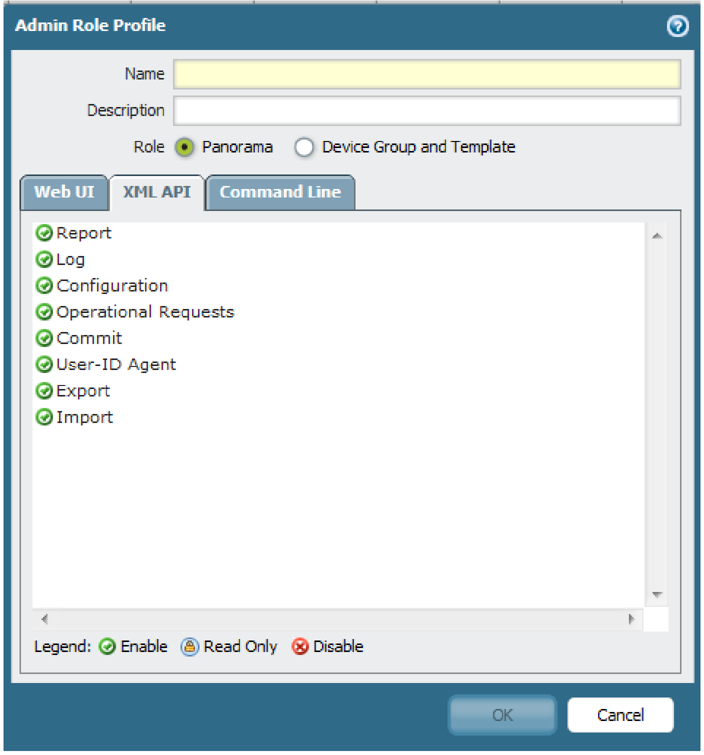

Name

Description

Role ● Panorama ○ Device Group and Template

✓ Report
✓ Log
✓ Configuration
✓ Operational Requests
✓ Commit
✓ User-ID Agent
✓ Export
✓ Import

Legend: ✓ Enable ◉ Read Only ⊗ Disable

6 Error Codes

The API response XML contains a status field and additionally an error field. The different error codes returned by the API in the error field are listed in the table below.

Error code	Name	Description
400	Bad request	Returned when a required parameter is missing, an illegal parameter value is used.
403	Forbidden	Returned for authentication or authorization errors including invalid key, insufficient admin access rights.
1	Unknown command	The specific config or operational command is not recognized.
2-5	Internal errors	Check with technical support when seeing these errors.
6	Bad Xpath	The xpath specified in one or more attributes of the command is invalid. Check the API browser for proper xpath values.
7	Object not present	Object specified by the xpath is not present. E.g. entry[@name='value'] where no object with name 'value' is present.
8	Object not unique	For commands that operate on a single object, the specified object is not unique.

9	Internal error	Check with technical support when seeing these errors.
10	Reference count not zero	Object cannot be deleted as there are other objects that refer to it. E.g. address object still in use in policy.
11	Internal error	Check with technical support when seeing these errors.
12	Invalid object	Xpath or element values provided are not complete.
13	Operation failed	A descriptive error message is returned in the response.
14	Operation not possible	Operation is not possible. E.g. moving a rule up one position when it is already at the top.
15	Operation denied	E.g. Admin not allowed to delete own account, Running a command that is not allowed on a passive device.
16	Unauthorized	The API role does not have access rights to run this query.
17	Invalid command	Invalid command or parameters.
18	Malformed command	The XML is malformed.
19-20	Success	Command completed successfully.
21	Internal error	Check with technical support when seeing these errors.
22	Session timed out	The session for this query timed out.

7 API Browser

The API browser is available at **http(s)://hostname/api.** You need to be logged in to the device's WebUI to be able to view the API browser.

You can use API browser to navigate different API requests that are available for use. For configuration commands, you can navigate to any path and view the corresponding xpath and API URL on the browser.

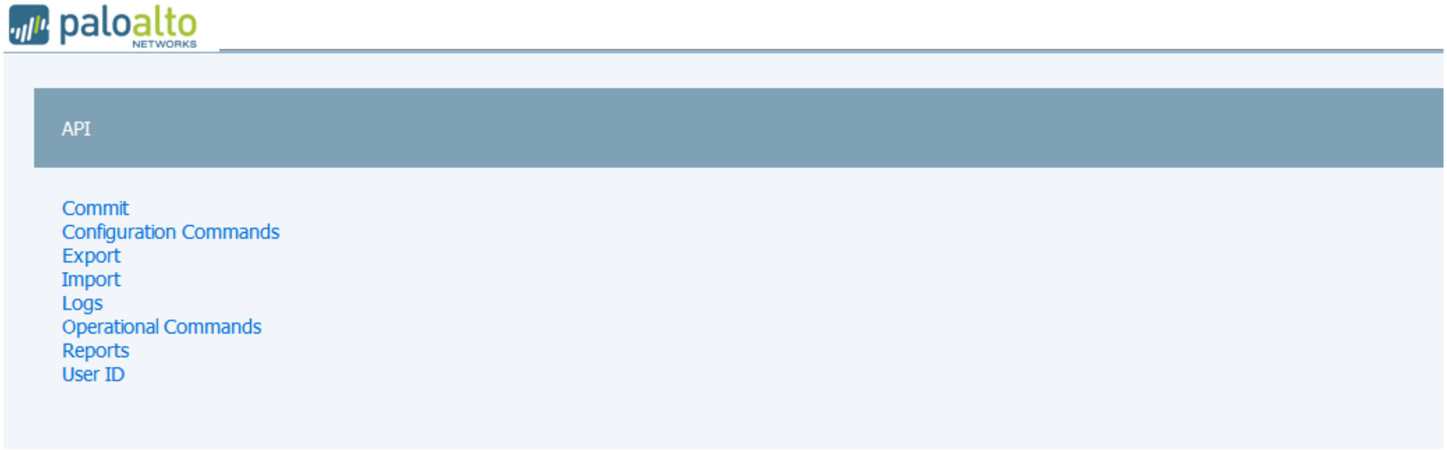

API

- Commit
- Configuration Commands
- Export
- Import
- Logs
- Operational Commands
- Reports
- User ID

For Configuration commands, you can navigate to a specific command to see its xpath.

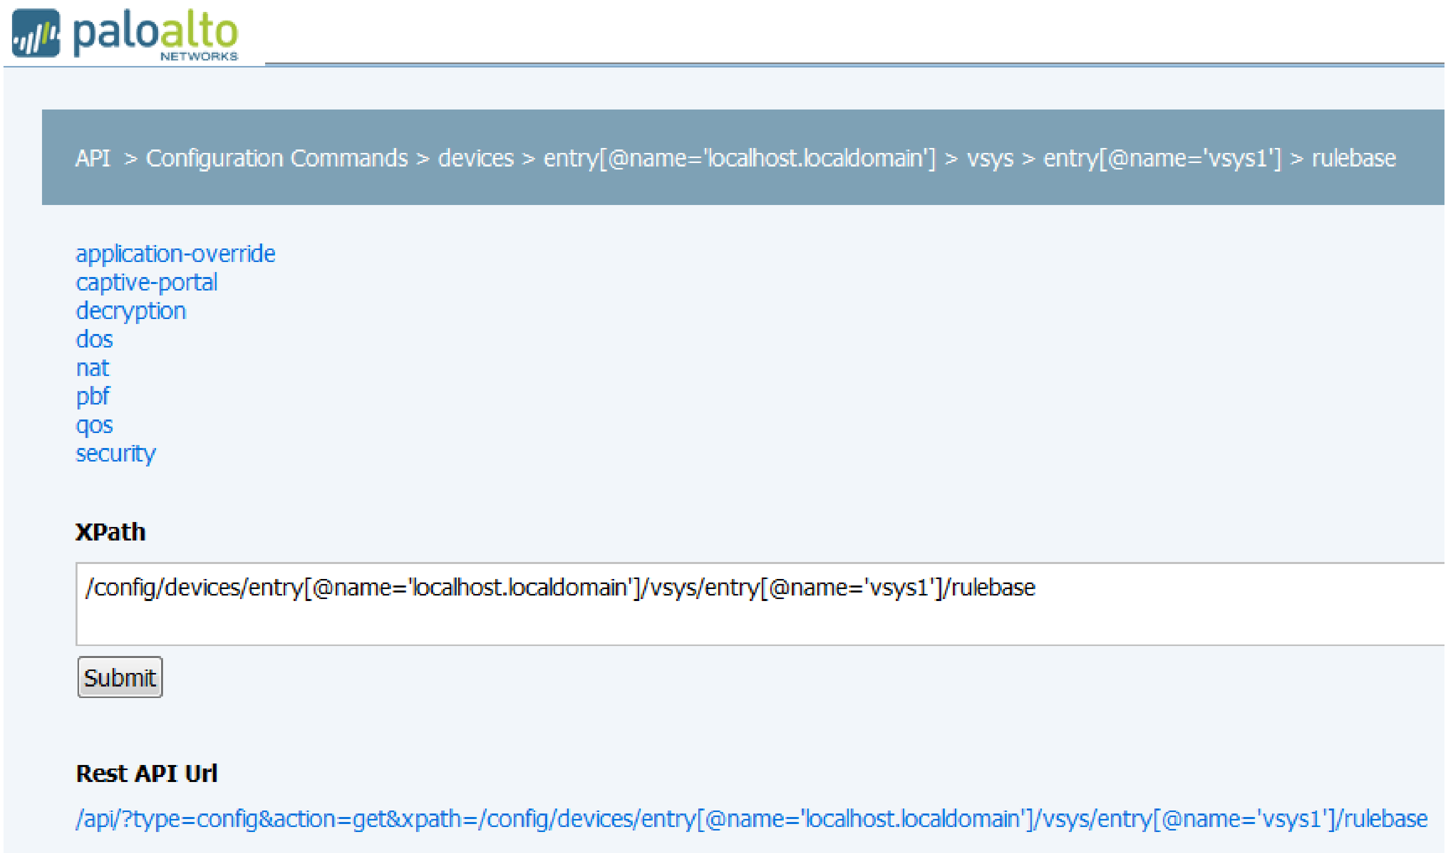

API > Configuration Commands > devices > entry[@name='localhost.localdomain'] > vsys > entry[@name='vsys1'] > rulebase

application-override
captive-portal
decryption
dos
nat
pbf
qos
security

XPath

```
/config/devices/entry[@name='localhost.localdomain']/vsys/entry[@name='vsys1']/rulebase
```

Rest API Url

```
/api/?type=config&action=get&xpath=/config/devices/entry[@name='localhost.localdomain']/vsys/entry[@name='vsys1']/rulebase
```

For Operational commands and Commit commands, you can navigate to a specific command to see the xml body to use for the *cmd* parameter.

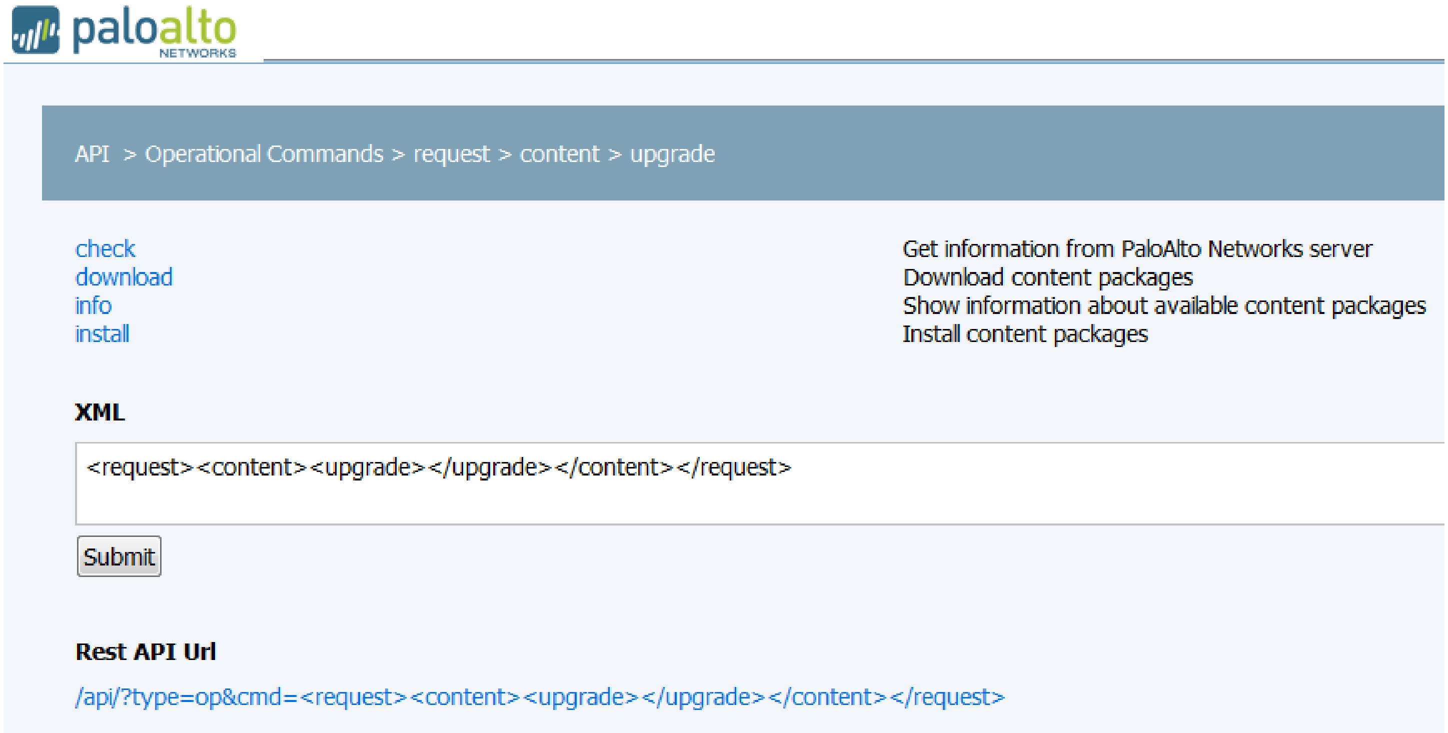

API > Operational Commands > request > content > upgrade

check	Get information from PaloAlto Networks server
download	Download content packages
info	Show information about available content packages
install	Install content packages

XML

```
<request><content><upgrade></upgrade></content></request>
```

Rest API Url

```
/api/?type=op&cmd=<request><content><upgrade></upgrade></content></request>
```

For reports, you can view the report names for all the supported dynamic and predefined reports.

8 Frequently Asked Questions

1 *How do I discover the xpath for the configuration object I am interested in?*
Use the API browser at http(s)://hostname/api to see all the available configuration commands along with their xpaths shown on the bottom of the screen. Alternatively, you can use the XML response for API request to show the entire running config, to navigate and discover the xpath for any element in the config.
http(s)://hostname/api/?type=config&action=show

2 *How do I build an xpath when there are multiple entries in a node in the config path to the element I am interested in?*
When there are multiple entries in any node in the path, you can specify the entry you are interested via the name of the entry, like so entry[@name='value']. For instance, the xpath to the address objects in vsys1 is /config/devices/entry/vsys/entry[@name='vsys1']/address

3 *How do I build the the xml body for the cmd parameter to be used in Operational and Commit commands?*
Use the API browser to navigate to a specific command and view the xml body to be used with the cmd parameter.

4 *Do I need to use URL/percent-encoding?*
You need to use URL encoding when using tools like cURL or wget. When using the browser, most browsers automatically do the URL encoding.

5 *What if my API request is too long?*
When the API request is 2K or longer, you should use HTTP POST instead of GET to avoid errors from the webserver. If you are using scripts and not a browser, you can use cURL or wget. Examples usages are shown below. Refer to their respective man pages for additional usage information.

Wget provides the --post-data and the --post-file options to do a HTTP POST.
> wget --post-data "*query-parameters*" http(s)://hostname/api/?query-parameters
> wget --post-file *input-filename* http(s)://hostname/api/?more-query-parameters, where the *input-filename* contains additional query paramaters for the API request.

Curl provides the --data options to do a HTTP POST.
> curl --data "*query-parameters*" http(s)://hostname/api/?*more-query-parameters*
> curl --data @input-filename http(s)://hostname/api/?*more-query-parameters*, where input-filename contains additional query parameters for the API request.

6 *How do I retrieve Panorama-pushed shared configuration from a firewall device?*
Use the Config Get API with xpath=/config/panorama. One example of this is if you want to retrieve pre- and post-rules from security policy.

7 *What are the xpaths and API queries for some sample configuration objects on the Firewall and Panorama?*
Creating a new URL filtering profile with a block action for www.badsite.com:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/profiles/url-filtering/entry[@name='xml test']&element=<description>xml api test</description><dynamic-url>yes</dynamic-url><action>block</action><block-list><member>www.badsite.com</member></block-list>

Adding a url to block list in an existing url profile:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/profiles/url-filtering/entry[@name='xml test']/block-list&element=<member>www.badsite.com</member>

Creating a new custom URL category:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/profiles/custom-url-category/entry[@name='xmltest urlcat']&element=<description>testing xml api</description><list><member>www.somesite.com</member></list>

Adding a URL to a custom URL category:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/profiles/custom-url-category/entry[@name='xmltest urlcat']/list&element=<list><member>www.somesite.com</member></list>

Adding an address object:
http(s)://hostname/api/?type=config&action=set&xpath=/config/devices/entry/vsys/entry[@name='vsys1']/address/entry[@name='xmltest addr']&element=<ip-netmask>1.2.3.4/32</ip-netmask><description>xml testing</description>

8 *How to pull Application and Threat Content information from the Firewall?*
Get a list of all the applications:
http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/application
Get a list of all the vulnerabilities:
http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/threats/vulnerability
Get information on a specific vulnerability by Threat-ID:
http(s)://hostname/api/?type=config&action=get&xpath=/config/predefined/threats/vulnerability/entry[@name='30003']

Revision History

Date	Revision	Comment
June 21, 2013	B	Added information on running an on-demand report on Panorama.
November 5, 2012	A	First release of the document.

Exhibit 19

paloalto
NETWORKS

Architecting User Identification (User-ID) Deployments

Strategies and Tactics guide
PAN-OS 4.1+

•

Contents

Section 1: User Identification software components

There are two major jobs for the User Identification (User-ID) process on the Palo Alto Networks firewall.

1) User to Group mapping: This process learns group names then the users that are members of the groups. It allows the firewall to write policy and create reports based on groups rather then individual users.
2) User to IP mapping: This process associated a user name with a specific IP. A given user can be associated with many IP addresses, but a single IP address can only be associated with a single user. The only exception to this is the Terminal Server Agent which is used for multi user Terminal Servers.

The Palo Alto Networks User Identification process consists of three separate components. The **Palo Alto Networks User-ID Agent**, the **Palo Alto Networks Terminal Server Agent** and the firewall running **PAN-OS.** Each component has specific functions and will require different topologies to function optimally.

- Palo Alto Networks User-ID Agent: This is a service that can be installed on any domain member system or run on the PAN-OS 5.0 firewall. It is responsible for 4 mechanisms that map users to their corresponding IP addresses.
  - Server Monitoring – Active Directory Domain Controllers, Microsoft Exchange Servers and Novell eDirectory servers.
  - Session Monitoring – Windows servers.
  - Client Probing – Windows systems thru either the WMI or NetBIOS.
  - User ID API – Extensible interface to import user data for other external sources then the ones mentioned above. This can include user defined scripts as well as partner integrations (such as Aruba Clearpass and Splunk)
- Palo Alto Networks Terminal Server Agent: This is a Windows service designed to be installed on a Windows or Citrix terminal Server to map users to IP + Source Port tuple. This provides tight correlation of terminal server traffic to terminal server user. It is required on any terminal server supporting users.
- The firewall software performs three specific functions in the User Identification process.
  - User to IP Mapping via Captive Portal – Both NTLM and Web Form methods are supported.
  - User to IP mapping using the Global Protect client in either internal or external modes.
  - User to Group mapping – Using LDAP the firewall will build a list of groups and the associated users for use in both policy and reporting. Note that the only component in a 4.1 environment that is required for user to group mapping is the PAN-OS firewall.

External systems referenced by User Identification

Configuration and design of User Identification can be complicated due to the number of external systems that may be part of the network infrastructure. Common external systems that are referenced in this document include:

- Windows Active Directory Domain Controllers
- Windows Exchange Server Client Access Servers
- Novell eDirectory Servers
- Other more generic LDAP implementations (OpenLDAP being the most common)

Section 2: Designing User to Group Mapping

Active Directory

By far the most common directory in current deployment is Microsoft's Active Directory. This is a LDAP based directory that has an extensive interface layered over it. This has the effect of masking much of the underlying LDAP structure from the casual administrator. The following is a list of common terms that play a part in the design of user to group mapping within Active Directory (AD)

- Active Directory Forest: This is the term applied to a full LDAP tree. A forest is a set of domains that share a common LDAP schema and have implied trust relationships running throughout. It is most common to have a single forest in a single corporate environment. Mergers and acquisitions may create scenarios where there are multiple forests in a single customer environment. Domains in a forest may have different names but they share common configuration and schema. The forest derives its name from the first domain created in the forest. For example if the first domain created was corp.com the forest would be referred to as the corp.com forest.
- Active Directory Tree: This is a Microsoft term for a set of domains in a single forest that share a common naming space. For example asia.corp.com and corp.com would be in the same tree. Domains in the same tree are by definition in the same forest. This term is immaterial to User ID design.
- Active Directory Domain: This is the smallest unit of LDAP replication in the Microsoft Active Directory. All AD deployments must consist of at least one Domain. A domain is represented by a single common name within the forest. For example asia.corp.com could be the name of a specific domain in the corp.com forest.
- Active Directory Domain Controller: A Domain Controller (DC) is a server that contains part of the Active Directory LDAP database. All DC's in a single domain have identical databases. This database contains all attributes for all of the objects in the Domain. By default each of the DC's is able to make changes to this database and then replicate those changes to other DC's in their domain. Standard DC's contain no information regarding objects in other domains from the forest.
- Global Catalogue Server: The Global Catalogue Server (GC) is standard Domain Controller for one of the domains in the forest that carries an additional LDAP database replica. This additional replica contains pointers for all of the objects in all of the domains in the forest, but does not contain all of their attributes. This GC LDAP database runs on a different port then the server traditional domain database. Specifically this database contains useful information concerning Universal Group membership.
- NetBIOS Name: While Active Directory has supported DNS style names since 2003 it is still common to see users and domains using the 15 max character NetBIOS name. A user named John Smith in the corp.com domain could be identified with a name such as **jsmith@asia.corp.com** or they could be referred to as asia/jsmith. In this case the domain asia.corp.com is referred to using its NetBIOS name. It is important to note that a domains NetBIOS name is not always the left most portion of its fully qualified name. For example a domain named johnsonbrothers.com may have a NetBIOS name of johnson.

-

Single Active Directory Domain

In a single AD domain environment user to group enumeration is simple. The following objects must be defined on the firewall:

1) LDAP authentication server
2) User Identification Group Mapping Settings

In a single domain, all Domain Controllers will have the identical information. As a result the firewall should be configured to connect to the nearest or best connected domain controller to gather the user and group data. Additional Domain Controllers can be added to provide fault tolerance but should be added based on their proximity to the firewall.

Single Active Directory Domain

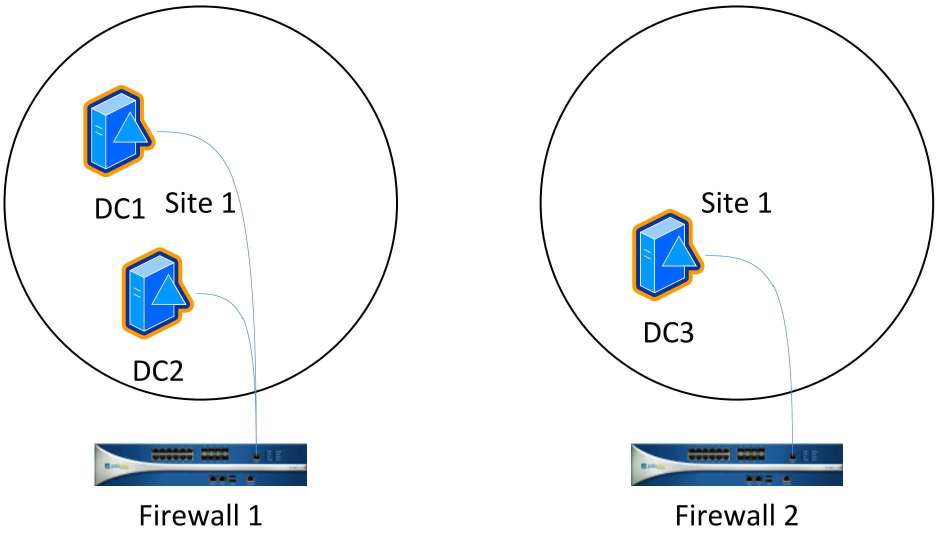

Multiple Active Directory Domains, Single Forest

In the case of a multiple AD Domain environment the firewall will need a separate LDAP server configuration for each domain. If Universal Groups are to play a significant role in firewall policy then an additional connection to a Global Catalogue server should be required as well. As in section 2.2 the firewall should connect to local instances of the domain controllers when possible. If no local instances are available for one of the domains then they firewall can be configured to connect to a remote domain controller with the update interval lengthened to reduce traffic across the WAN.

Required objects are:

1) LDAP authentication server for each domain and optionally for the Global Catalogue.
2) User Identification Group Mapping setting for each LDAP server.

Multiple Active Directory Domain

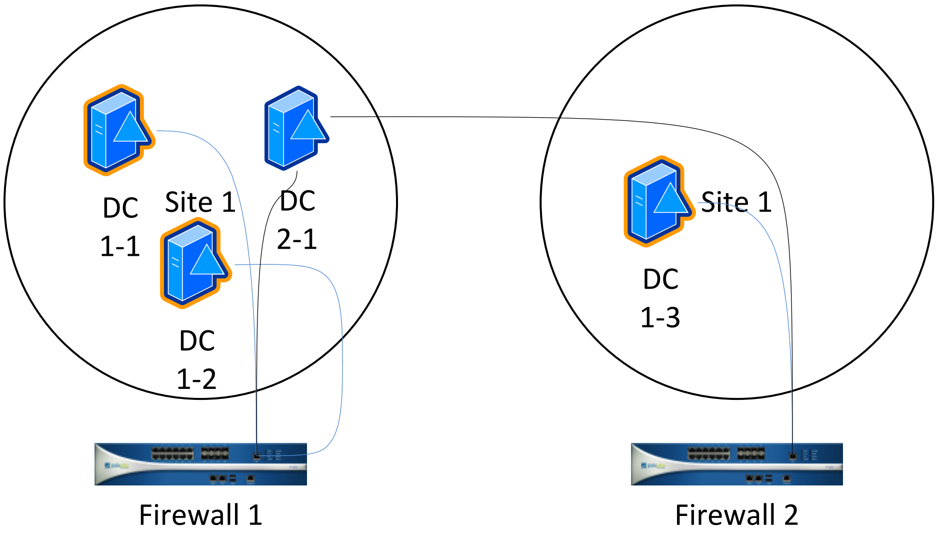

Multiple Active Directory Forest

From the point of view of the firewall and group enumeration, a multiple forest environment is identical to the multiple domain environment. The additional concern with multiple forests is that there is no guarantee of unique names with respect to the other forest. Care will need to be taken that two identically named groups / users from two identically named domains are not encountered.

Section 3: Designing User Agent Topology for User to IP Mapping

The User-Identification Agent

To successfully deploy an identity aware security solution using Palo Alto Networks firewalls users must be mapped to IP addresses in real time. In smaller and simpler networks the deployment of a software or hardware user agent in a central location will suffice to map the majority of users. When planning agent placement it is important to consider the traffic generated by the agent. For the software agent the traffic generated is as follows:

- Log monitoring traffic: This traffic occurs between the agent and an Active Directory Domain Controller or Exchange CAS. The content of this traffic is the entirety of the servers Security Log. Based on the “Security Log Monitor Frequency” value configured on the agent this traffic consists of an authenticated TCP based WMI connection to the server that fetches the new logs since the last check.
- Open server session traffic: This traffic occurs between the agent and an Active Directory Domain Controller or Exchange CAS. The content of this traffic is the entirety of the servers’ current session table for file and print shares. Based on the “Server Session Read Frequency” value configured on the agent this traffic consists of an authenticated TCP based WMI connection to the server that fetches the current session table.
- Probing (WMI and NetBIOS) traffic: If probing is enabled the agent will make a connection using either the WMI or NetBIOS to each known IP address discovered through log or session monitoring. This connection will be used to verify that the last known user is still the current user. The list of known IP addresses is cycled through once each period defined by the WMI/NetBIOS Probing Interval as configured on the agent.
- Firewall update traffic: when a user mapping is learned or updated the agent pushes this data to the firewall. In addition the agent refreshes all known mappings to the firewall every hour. This traffic consists of just the username and corresponding IP address as well as a time stamp.

Traffic between Domain Controller and Agent, or Exchange Server and Agent cannot be optimized. The amount of traffic is determined by the amount of logging activity on the target server. As a result it is a best practice to place the agent closer to the monitored servers when bandwidth is a concern. Traffic between the Agent and the Firewall consists of the bare minimum data required for user to IP mapping and is better suited to traverse more impacted links.

Basic agent placement theory – Log Reading

It is usually a best practice to place Agent systems near by the Domain controllers they will monitor. “Near” being a relative term with respect to networking. In the simplest design, an agent is placed in each physical site that is separated by a WAN link and that contains DC’s or Exchange servers. Since users could theoretically authenticate to any DC in the environment and since the security logs are not replicated between Domain Controllers, all DC’s in the enterprise must be monitored. Domain Controller / Exchange log monitoring is the lowest overhead option the agent provides and should always be used as the base method for user to IP mapping.

Each firewall in the deployment should receive updates from every agent. By placing agents across impacted WAN links we minimize the User Identification traffic over the links. The remote agent can query local DC's and then send he summarized data back to the firewall.

Basic agent placement theory – Probing

Agent probing can be used to verify the known user is still at an IP address that was learned from another method and for learning the user at an unknown IP address referenced by the firewall. Probing generates network traffic based on the total number of learned IP addresses of the network. Probing is most useful in networks with a high turnover of user to IP address. The probes can time out or update an IP mapping before the cache timer is hit. This can give a tighter correlation of current user to IP status.
WMI probing is always preferred to the legacy NetBIOS option. (The hardware PAN-OS agent does not support NetBIOS)
Advantages of WMI probing are:

1. WMI probes are authenticated. The agent will use its account to authenticate itself to each end point.
2. WMI is a more reliable method of gathering user data then NetBIOS
3. WMI is more likely to be allowed within the network then NetBIOS

Probing is best used on a well-connected network. It is not appropriate for WAN or congested network segments. In cases where the IP to user mappings are relatively static, probing is most likely unneeded.

The diagram below shows a simple topology using a network with 2 sites and 2 domains. Each site has an agent that is responsible for monitoring the Domain Controllers in that site. The single centralized firewall receives updates from both agents. The total number of agents in this design is 2.

Simple Agent Topology

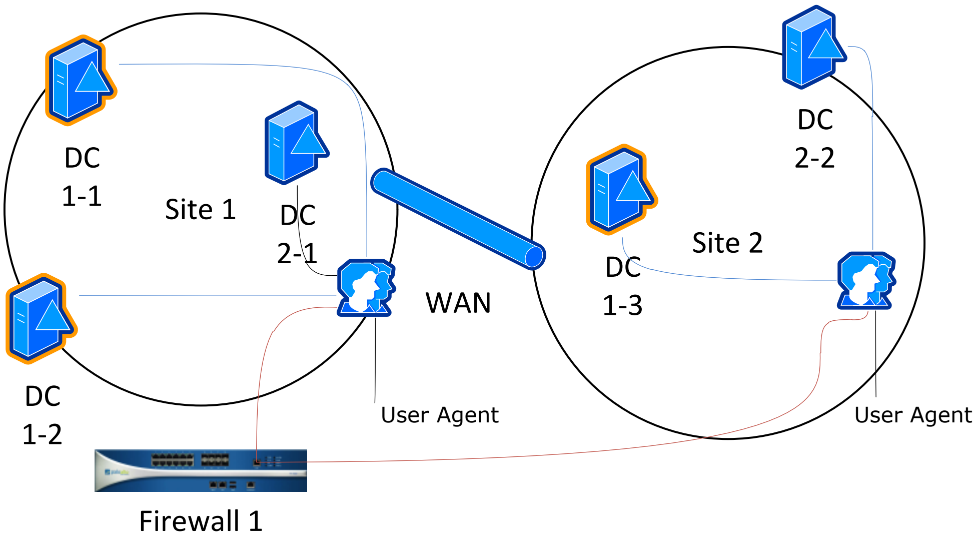

Section 4: Captive Portal

Captive portal provides an active method to authenticate an unknown user. It will only be triggered when traffic from an unknown IP matches a captive portal policy. Since it is only invoked for this unknown traffic it provides a low cost option for identifying any IP not covered in the log or session monitoring. Most designs can benefit from the addition of captive portal to provide mapping for users that slip through the cracks of the other methods. Captive portal is only effective on web traffic. For firewalls deployed in a region of the network that does not encounter web traffic often, this method is less useful.

Captive Portal modes

Captive portal can function in 2 modes. Depending on the mode supported by clients the design consideration for captive portal are different.

- **NTLM Authentication:** If the client systems are running Windows and they are using Firefox or Internet Explorer as a web browser then NTLM authentication can be used. NTLM provides a transparent authentication method that will not impact the user experience. In environments where this is an option it is a best practice to enable it. For environments that are Microsoft and the browser usage can be controlled this method provides a seamless backup to the log monitoring. It will not work in environments where the users authenticate to local systems with a different account then the one used in the Active Directory.
- **Web Form Authentication:** In all environments other than described above, captive portal will use web form. Web form can also be enabled for all use cases if desired. An advantage to web form is that the user can be authenticated using any of the available authentication sources. A disadvantage is that this method interrupts the users work to prompt for credentials. This method is ideal for kiosk systems where the local user is not a user from the enterprise directory and where many different users may utilize the same system without performing any network log on.

PAN-OS 5.0 and later firewalls and be configured to redistribute this data to other firewalls in the environment.

Section 5: Terminal Services Agent

For Windows and Citrix servers a special agent must be used. These systems multiplex users behind a single IP address. Traditional User ID methods cannot address this type of user traffic. The TS Agent is a small footprint agent that must be installed on each terminal server. The agent will control the source ports allocated to each user process and report this to the firewall. It is the only user ID component required for terminal servers.

From the point of view of design there is only one way to deploy the TS Agent. It must be installed on every Terminal Server and then added to each firewall that may encounter traffic from the terminal server and will require user data.

Section 6: Global Protect

The Global Protect client software plays a role in the Palo Alto Networks remote access, mobility and User Identification solutions. With respect to User ID Global Protect acts as a trusted client process that can report the user and IP of the endpoint to PAN-OS firewalls acting as gateways.

If an environment requires 100% correlation of user to IP at all times then the Global Protect client with Internal Gateways is the best solution. This does not need to be pushed out globally. It is possible that only a sub set of the client systems require this tight correlation. The rest of the user base may be serviced by log monitoring and Captive Portal.

Section 7: Advanced Designs

When dealing with larger or more complex networks, additional techniques may be employed to cover user to IP mapping. The strategies covered here can be mixed and matched to address any customer environment. It is important to realize that in most customer networks, multiple strategies may need to be employed to reach an acceptable level of user to IP mapping.

Techniques

The following techniques will be covered in this guide:

1) Central deployment of a hardware PAN-OS User agent for distributed environments
2) Deployment of software agents for targeted segments within the network
3) Use of Microsoft Log forwarding for highly distributed networks
4) Use of smaller (PA-200) firewalls as dedicated User ID appliances

Scenarios

The following broad descriptions represent the most common customer networks that require more sophisticated User Identification solutions:

1) Highly distributed, low density Domain Controllers. – Networks containing many Domain Controllers but having no clear hub sites. Sets of one or two DC's connected to other sites over a WAN. No logical location for centralized agent placement. Example: A regional financial institution with DC's at each branch.
2) Multiple high latency / low bandwidth / heavily subscribed links. – Networks containing sites with Domain Controllers that are poorly connected to central locations. Example: Oil platform with satellite link
3) Large WiFi segments with high user turnover. – Segments where users are ephemeral and the turnover is high. Login events are not frequent enough to map users as they move through the network, Example – Large campus WiFi
4) Distributed sites with agent HA requirements – Networks with a number of remote sites (10+) that require agent High Availability. This requires a large number of software agents to be configured.
5) User ID and multiple VSYS – User agents are not shared between virtual systems. A network with 10 user agents and 5 VSYS would effectively have 50 agents configured. The hard limit of 100 agents per firewall limits the design options for large VSYS deployments.

Hardware PAN-OS Agents

While both the stand alone software User ID agent and the "agentless" PAN-OS processes perform the same basic tasks they use different underlying protocols. This difference makes each one more appropriate for different environments. The software agent uses MS RPC to query the Domain Controller and Exchange Server logs. This method requires the full log to be transferred to the agent where it is then filtered for the required events. The hardware integrated agent uses the

WMIC library and only transfers the required log events to the agent process.

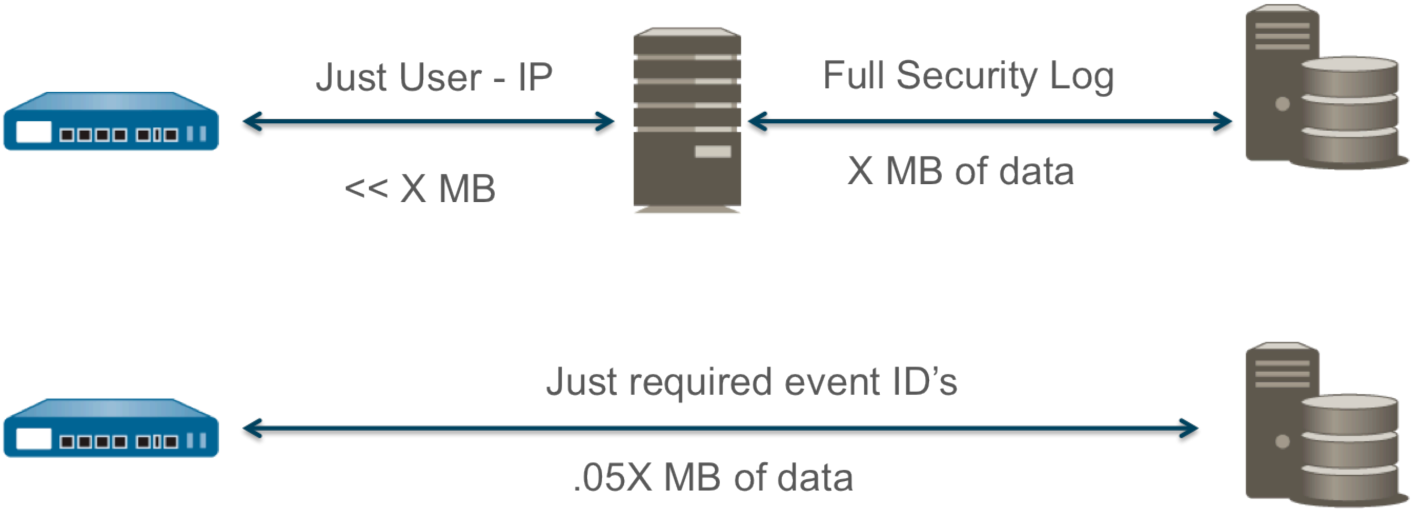

As a result the hardware agent is appropriate for reading remote Domain controllers where the software agent is appropriate for reading local Domain Controllers. The drawbacks to the hardware agent are the following:

1) Resources for the agent process come from the Management Plane. Significant User ID activity can impact other management plane features such as reporting and management.
2) There is no way to increase the resources for the hardware agent as the User ID environment grows.
3) The WMIC can have a higher impact on the target Domain Controllers CPU. This is mostly noticeable on 32 bit low RAM servers.

This technique can allow a significant reduction in the total number of agents required by allowing the agent process to sit in a central location rather then in multiple remote sites.

Performance of the hardware agent is determined by the firewall model performing the service. The following table shows preferred numbers for each platform. If there are significant other requirements on the management plane the number of DC's should be reduced.

Platform	DC's supported
4000, 2000, 500	10
200	25
3000, 5000	100

Use of the hardware agent is suggested for the following scenarios:
Highly distributed, low density Domain Controllers.
Multiple high latency / low bandwidth / heavily subscribed links.
User ID and multiple VSYS

Targeted / Multiple Software Agents

All settings on the User Agent are global. There is no facility for different agent setting values based on the network range. In some cases different network segments would benefit from different agent settings. To accomplish this multiple agents will need to be installed and their "Include / Exclude" network lists need to be configured so that they divide the customer environment into the correct segments.

. There are a number of valid reasons to install more than one agent in a single site.

1. **Fault Tolerance:** Multiple agents monitoring the same domain controllers provide redundancy for the firewall should one of the agents fail. The multiple agents do not need to be aware of each other. The agents singularly send all known IP to user mappings to the firewall. The firewall will normalize the multiple feeds and remove any duplicate data.
2. **Requirement of different agent settings for different network segments:** Agent global settings such as the enablement of probing or cache time out may not be ideal for all subnets at a given site. For example a common wireless network at a main site may have a very short DHCP lease time and hence a fast turnover of users to IP addresses. The rest of the network may be wired and have week long lease times. In this case a cache time out of 120 minutes for the wired network would be too long for the wireless. In addition the wireless network may benefit from WMI probing while it would be a significant traffic increase to probe the full wired network with very little benefit. Since both settings are global to the agent the only way to provide probing and a 15 minute time out to the wireless network while disabling probing and assigning a 120 minute time out to the wired network is to install two agents. Each agent would have their Include / Exclude list configured to only cover their portion of the sites IP address space. Each agent would be configured to monitor all of the same DC's and Exchange servers.
3. **Load Balancing:** In a large data center there could be a very high concentration of Domain Controllers or Exchange servers with very high levels of logging. The available monitored servers could be split up between multiple agents to minimize the amount of logs each agent was required to process.

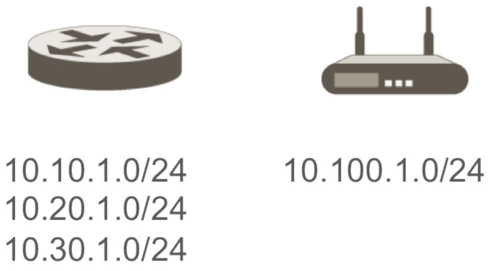

Agent 1	
Probing Setting	Off
Include /Exclude List of configured Networks	10.10.1.0/24,10.20.1.0/24, 10.30.1.0/24

Agent 2	
Probing Setting	On
Include /Exclude List of configured Networks	10.100.1.0/24

The agent values most likely to be targeted in this way are the WMI Probing Setting and the Cache time out.

Use of targeted software agents is suggested for the following scenarios:
Large WiFi segments with high user turnover.

Microsoft Log Forwarding

Microsoft servers support a publish / subscribe service for forwarding event logs from one server to another. We can use this service to the required security logs from remote Domain Controllers to a server more convenient for the software or hardware agent to monitor. This is a built in feature from Microsoft and does not require any additional software. This can be used to effectively reduce the total number of Domain Controllers that need to be monitored. This service will batch the logs over 30 seconds for efficient transport. It does introduce this additional latency into the user ID process. MS Log Forwarding is supported on all versions of Windows Server.

Advantages of MS Log forwarding:

1) Reduce the number of servers that User Agents need to monitor.
2) Publish logs from remote or poorly connected sights.
3) Only the required logs need to be sent to the central server.
4) This configuration can be made standard on new DC's mitigating the need to keep the agent up to date on all existing DC's in production.
5) The agent only needs rights on the subscribing server, which does not need to be a Domain Controller.

-
- Disadvantage of MS Log Forwarding
  1) Can introduce up to an additional 30 seconds of latency into the process.
  2) Server session reading is no longer useful as the agent is not connecting to the actual Domain Controller.
-

Use of MS Log Forwarding is a possible recommendation for the following scenarios:
2. Highly distributed, low density Domain Controllers.
3. Multiple high latency / low bandwidth / heavily subscribed links.
Distributed sites with agent HA requirements.
User ID and multiple VSYS

Dedicated Hardware Agents

As an extension of the PAN-OS on board agent solution, dedicated platforms can be used to provide Agent services. Most commonly we deploy PA-200's in this manner but VM series systems could be used as well.

Dedicating hardware to this process has the following advantages:

1) Platforms can be deployed in HA pairs giving fault tolerance without adding additional agents to the configuration. This can drastically reduce the total number of agents configured on the traffic forwarding firewalls.

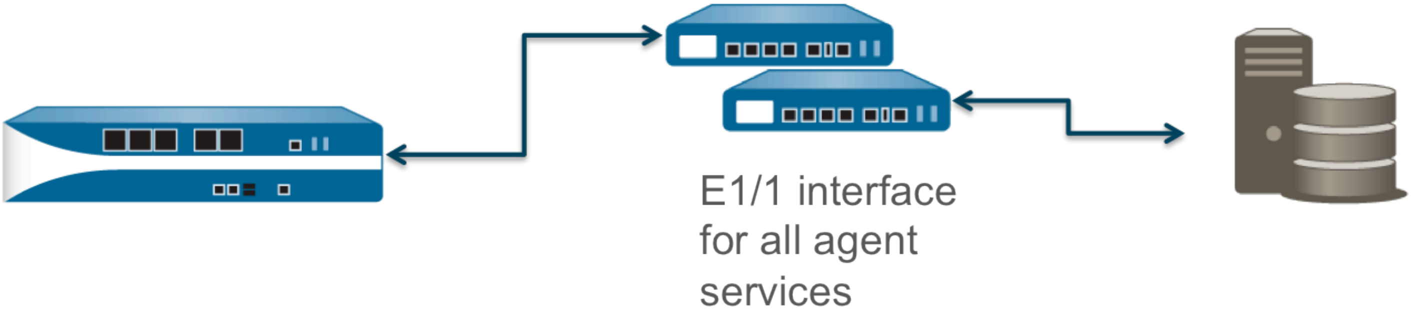

2) The management plane CPU and RAM are used fully for the User ID service allowing the platform to monitor the maximum number of DC's recommended.
3) Upgrades and configuration of the hardware agents can be done centrally though Panorama.
4) In some cases the removal of the Windows server requirement would be seen as beneficial.

Use of Dedicated Hardware Agents is recommended for the following scenarios:
Distributed sites with agent HA requirements.
User ID and multiple VSYS

Revision History

Date	Revision	Comment
9/11/13	.9	Draft
9/13/13	.91	Re-done with more coverage
10/1/13	1	First published version

Exhibit 20

Traps: Advanced Endpoint Protection

TRAPS:

- Prevents all vulnerability exploits
- Prevents all malware-driven attacks
- Provides Immediate forensics of prevented attacks
- Is scalable, lightweight and user friendly
- Integrates with the network and cloud security

Palo Alto Networks® Traps provides Advanced Endpoint Protection that prevents sophisticated vulnerability exploits and malware-driven attacks. Traps accomplishes this through a highly scalable, lightweight agent that uses an innovative new approach for defeating attacks without requiring any prior knowledge of the threat itself. By doing so, Traps provides organizations with a powerful tool for protecting endpoints from virtually every targeted attack.

Palo Alto Networks Traps takes a unique approach to endpoint security, designed to provide complete security protection for the endpoint, including the prevention of both conventional attacks as well as advanced and targeted attacks that traditional solutions cannot prevent.

Instead of looking to identify the millions of individual attacks themselves, or detect malicious behavior that may be undetectable, Traps focuses on the core techniques that every attacker must link together in order to execute their attack. By setting up a series of exploit 'traps' into the process to mitigate these techniques, Traps can thwart the attack immediately before any malicious activity can successfully run.

This unique approach allows Traps to be agnostic to application, protecting all applications, including those developed by 3rd parties.

Exploit prevention

The actual process of exploiting a vulnerability on an endpoint requires execution of multiple advanced techniques operating in sequence. For example, in a typical attack the attacker will attempt to gain control of a system by first attempting to corrupt or bypass memory allocation or handlers. By using memory-corruption techniques such as buffer overflow or heap corruption the hacker can utilize weaknesses or vulnerabilities within the target software to execute their specific code. Once an attacker is able to execute custom code, he can download malware or completely control the system to his full advantage.

Regardless of the attack or its complexity—in order for the attack to be successful the attacker must execute a series exploit techniques in sequence. Some attacks may involve more steps, some may involve less, in all cases at least two or three techniques must be used in order to exploit the targeted endpoint.

How Exploit Prevention works

Traps employs a series of exploit prevention modules aimed at mitigating and blocking the different exploit techniques available to attackers. These modules operate like "traps", injected into the user processes and designed to trigger and block the attacker's exploit technique as soon as it's attempted. Whenever an application is opened Traps seamlessly injects prevention modules into the process as transparent, static "traps". Once a module is injected into the process, that process is then protected from any exploit. If an exploit attempt is made using one of the few available techniques, Traps will immediately block that

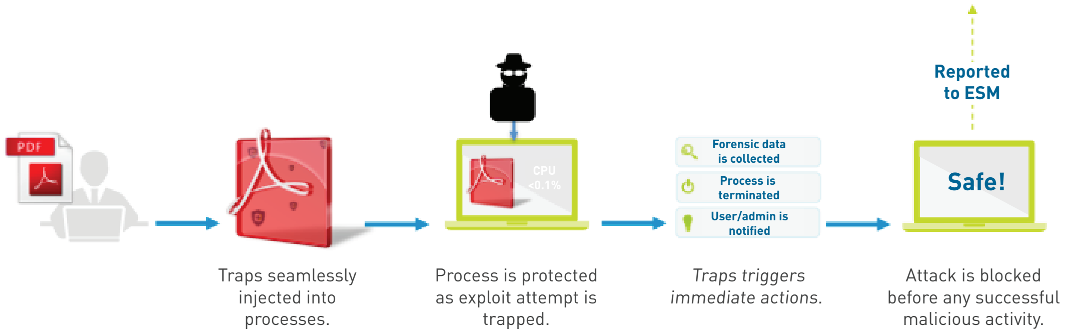
PDF
CPU <0.1%
Reported to ESM
Forensic data is collected
Process is terminated
User/admin is notified
Safe!
Traps seamlessly injected into processes.
Process is protected as exploit attempt is trapped.
Traps triggers immediate actions.
Attack is blocked before any successful malicious activity.

How it works: Exploit prevention.

technique, terminate the process, and notify both the user and the admin that an attack was prevented. In addition, Traps will collect detailed forensics and report that information to the Endpoint Security Manager (ESM). Due to the chain-like nature of an exploit, preventing just one technique in the chain is all that is needed in order to block the entire attack.

If no attempt is made it's business as usual for that user and process. Given the minimal resource utilization of Traps, there will be no user experience implications of the preventative measures that were deployed behind the scenes.

By focusing on the exploit techniques and not the attack itself, Traps can prevent the attack without prior knowledge of the vulnerability, regardless of patches in place, and without signatures or software updates. It's important to note that Traps isn't scanning or monitoring for malicious activity, so there's a massive scalability benefit to this approach as very little CPU and memory are used.

Traps exploitation prevention is designed to prevent attacks on program vulnerabilities based on memory corruption or logic flaws. Examples of attacks that Traps can prevent, include:

- Memory corruption
- Java code from running in browsers, under certain conditions
- Executables from spawning child processes, under certain conditions
- Dynamic-link library (DLL) hijacking (replacing a legitimate DLL with a malicious one of the same name)
- Hijacking program control flow
- Inserting malicious code as an exception handler

Malware Prevention

Malicious executable files, known as malware, are often disguised as or embedded in non-malicious files. They can harm computers by attempting to gain control, gather sensitive information, or disrupt the normal operations of the system.

While advanced attackers are increasingly exploiting software vulnerabilities, attacks are also advancing with unknown or manipulated Malware (Executable files) and because these types of attacks generally don't have known signatures, known strings or previously known behavior, traditional endpoint security approaches are unable to prevent them.

In order to effectively prevent the execution of malware on the endpoint, Traps employs the following three components of malware prevention:

1. **Policy-Based Restrictions:** Policy restrictions provide organizations with the ability to set up policies restricting specific execution scenarios, and not the whitelisting or blacklisting of specific files. The attack surface can be greatly reduced by simply controlling the source of file installation. When a user attempts to open the executable, Traps will evaluate the execution restriction rules that may apply. Examples of common policy based restrictions;
   - Running executables from certain folders
   - Running executables from external media
   - Processes spawning child processes
   - Java processes run from browsers
   - Running unsigned processes
   - Thread Injection
2. **Wildfire™ Inspection:** For execution of files that are not limited to the policy restrictions set in place, Traps Endpoint Security Manager will query the WildFire threat cloud with a hash, to determine if the file is malicious, benign, or unknown within the global threat community. If WildFire confirms that a file is known malware, Traps will prevent the file from executing and will notify the ESM.
3. **Malware Techniques Mitigation:** Similar to Exploit techniques,

attackers utilize common, and identifiable techniques when trying to deploy their malware. In the event that the file execution is not restricted by policy or has not been matched by hash to a known attack in the Wildfire threat cloud, Traps will implement technique-based mitigations that limit or block; child processes, Java processes initiated in web browsers, remote thread and process creation, and unsigned processes execution—in order to prevent the attack entirely from executing.

Forensics

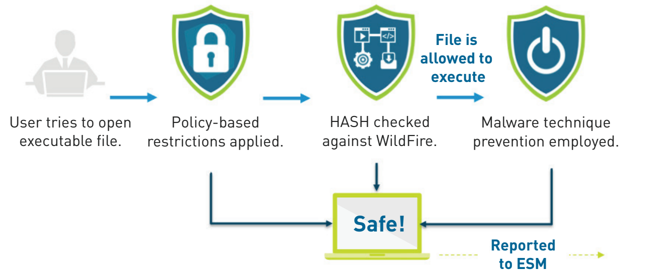
User tries to open executable file.
Policy-based restrictions applied.
HASH checked against WildFire.
File is allowed to execute
Malware technique prevention employed.
Safe!
Reported to ESM

How it works: Malware prevention.

Whenever Traps prevents an attack, real-time forensic details about the event will be collected about; the file, what occurred, the memory state when it was prevented, etc. and report the logged information to the Endpoint Security Manager (ESM). Despite the fact that the attack was prevented, there is still a great amount of intelligence that can be gathered. By capturing all the forensics of the attempted attack, organizations can apply proactive defenses to other endpoints that may not be protected.

Traps Architecture

Traps provides a 3-tier management structure consisting of the Endpoint Security Manager, Endpoint Connection Server, and endpoint agents. This model allows for massive horizontal scalability while still maintaining a centralized configuration and database for policies, forensics, etc.

Endpoint Security Manager

The Endpoint Security Manager provides an administrative dashboard for managing security events, endpoint health, and policy rules. The ESM also handles the communication to WildFire when hashes are sent for inspection. The ESMs all-in-one management center covers:

- Configuration management
- Logging and DB query
- Admin dashboard and security overview
- Forensics captures
- Integration configuration

The Endpoint Security Manager includes a centralized database that stores administrative information, security policy rules, endpoint history, and additional information about security events. The database is managed over the MS-SQL platform.

The Endpoint Security Manager can write logs to an external logging platform, such as security information and event management (SIEM), Service Organization Controls (SOCs), or syslog, in addition to storing its logs internally. Specifying an external logging platform allows an aggregated view of logs from all Endpoint Servers.

Endpoint Server

The Endpoint Server regularly distributes the security policy to all agents and manages all the information related to security events.

- **Traps Status** – Notifications and health pages in the Endpoint Security Manager display the status for each endpoint.
- **Notifications** – Traps agent sends notification messages about changes in the agent, such as the start or stop of a service, to the Endpoint Server.
- **Prevention reports** – Traps reports all of the information pertaining to an event, to the Endpoint Server in real-time.

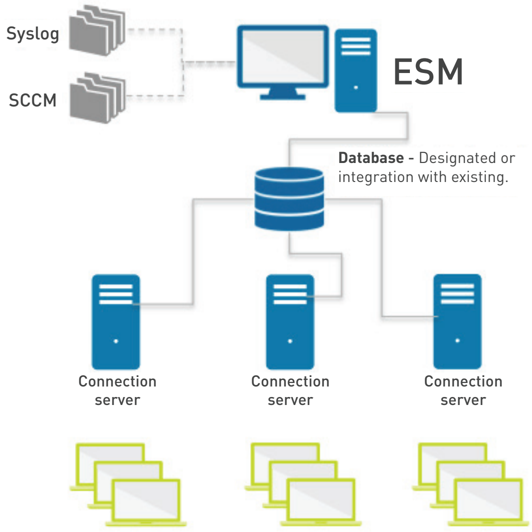

PCs, servers, VMs, VDI, Citrix session, thin client, embedded

Coverage & Platform Support

Traps protects unpatched systems, requires no hardware and is supported across any platform that runs Microsoft Window; desktops, servers, industrial control systems, terminals, VDI, VMs and embedded systems etc.

Traps currently supports the following Windows-based operating systems:

WORKSTATIONS

- Windows XP SP3
- Windows 7
- Windows 8.1
- Windows Vista SP1

SERVERS

- Windows Server 2003
- Windows Server 2008 (+R2)
- Windows Server 2012 (+R2)

Specifications

With the unique approach taken, Traps operates in a somewhat static capacity and doesn't scan for malicious activity our resource utilization is very low:

TRAPS AGENT:

- CPU – Average utilization of 0.1%
- Memory Consumption – 25 MB
- Disk Space – 15 MB

paloalto networks®
the enterprise security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main:	+1.408.753.4000
Sales:	+1.866.320.4788
Support:	+1.866.898.9087

www.paloaltonetworks.com

PAN_DS_TRAPS_100814

Exhibit 21

paloalto networks.

Home > Company > Investor Information > SEC Filings >

SEC FILINGS

SEC Filing Keyword Search (View search tips)

Year Filter

All Years

Groupings Filter (View SEC Groupings descriptions)

All Forms Search

View Section 16 Filings (3,4,5)

<< First | Previous | Next | Last >>

Filing Date	Form	Description	Filing Group	Downloads
10/30/14	DEFA14A	Additional proxy soliciting materials - definitive	Proxy Filings	
10/30/14	DEF 14A	Official notification to shareholders of matters to be brought to a vote ("Proxy")	Proxy Filings	
10/23/14	4	Statement of changes in beneficial ownership of securities	3,4,5	
10/23/14	4	Statement of changes in beneficial ownership of securities	3,4,5	
10/23/14	4	Statement of changes in beneficial ownership of securities	3,4,5	
10/15/14	4	Statement of changes in beneficial ownership of securities	3,4,5	
10/15/14	4	Statement of changes in beneficial ownership of securities	3,4,5	
		Statement of changes in beneficial ownership of		

- Overview
- Stock Information
- Corporate Governance
- SEC Filings
- Quarterly Financial Results
- News Releases
- Calendar of Events
- Information Request
- Investor FAQ

STOCK QUOTE

PANW (Common Stock)

Exchange	NYSE (US Dollar)
Price	**$106.71**
Change (%)	▲ 0.75 (0.71%)
Volume	1,000,950

Data as of 11/04/14 4:02 p.m. ET

Minimum 20 minute delay

Refresh quote

INVESTOR TOOLS

- Print Page
- E-mail Page
- RSS Feeds
- E-mail Alerts
- IR Contacts
- Share Page

Sales: 866.320.4788

Privacy Policy | Legal Notices | Site Index | Subscriptions |

Exhibit 22

Essentials 1: Firewall Installation, Configuration, & Management

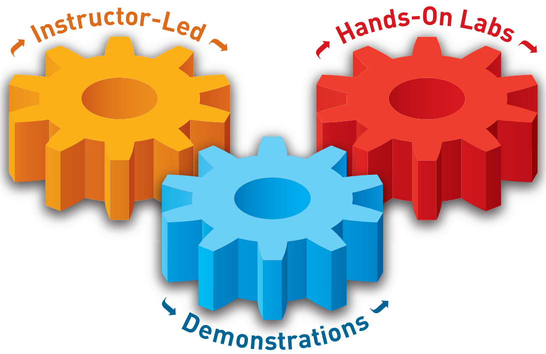

OVERVIEW

Successful completion of this three-day, instructor led course will enable the student to install, configure, and manage the entire line of Palo Alto Networks™ Next-Generation firewalls.

COURSE OBJECTIVES

Students attending this introductory-level class will gain an in-depth knowledge of how to install, configure, and manage their firewall, as well as configuration steps for the security, networking, threat prevention, logging, and reporting features of the Palo Alto Networks Operation System (PAN-OS).

SCOPE

- Course level: Introductory
- Course duration: 3 Days
- Course format: Combines lecture with hands-on labs
- Platform support: All Palo Alto Networks next-generation firewall models

TARGET AUDIENCE

- Security Engineers, Network Engineers, and Support staff

PREREQUISITES

Students must have a basic familiarity with networking concepts including routing, switching, and IP addressing. Students should also be familiar with basic port-based security concepts. Experience with other security technologies (IPS, proxy, and content filtering) is a plus.

COURSE OUTLINE:

DAY 1

Module 0 – Overview

Module 1 – Administration & Management
- GUI, CLI, and API
- Config Management
- PAN-OS & Software Update

Module 2 – Interface Configuration
- Layer 2, Layer 3, Virtual Wire, Tap
- Subinterfaces
- Security Zones

Module 3 – Layer 3
- Layer 3 Configurations
- Interface Management
- Service Routes
- DHCP
- Virtual Routers
- NAT (source and destination)
- IPv6 Overview

DAY 2 & 3

Module 4 – App-ID™
- App-ID Process
- Security Policy Configuration
- Policy Administration

Module 5 – Content-ID™
- Antivirus
- Anti-spyware
- Vulnerability
- URL Filtering
- File Blocking: WildFire™
- Zone Protection

Module 6 – Decryption
- SSL Inbound and Outbound

Module 7 – User-ID™
- User-ID Agent
- Enumerating Users
- Mapping Users to IP
- Users in Security Policy

Module 8 – VPN
- IPsec
- GlobalProtect Overview

Module 9 – High Availability
- Configuring Active/Passive

Module 10 – Panorama
- Device Groups & Templates
- Shared Policy
- Config Management
- Reporting and Log Collection

ORDERING INFORMATION:

PART NUMBER: PAN-EDU-201

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

Exhibit 23

Essentials 2: Extended Firewall Management

COURSE OUTLINE:

DAY 1

Module 0 – Overview

Module 1 – Administration & Management

- Password Management
- Certificate Management
- Log Forwarding

Module 2 – Interface Configuration

- VLAN Objects
- QoS

Module 3 – Layer 3

- NAT
- Policy Based Forwarding
- Routing Protocols (OSPF)

Module 4 – App-ID™

- Defining new Application Signatures
- Application Override

Module 5 – Content-ID™

- Custom Threat Signatures
- Data Filtering
- DoS Protection
- Botnet Report

DAY 2

Module 6 – User-ID™

- Captive Portal
- Terminal Server Agent
- XML API
- Dynamic Address Objects

Module 7 – VPN

- Overview
- Configuring the Portal, Gateway, and Agent
- Host Checks
- Logs

Module 8 – High Availability

- Configuring Active/Active HA

ORDERING INFORMATION:

PART NUMBER: PAN-EDU-205

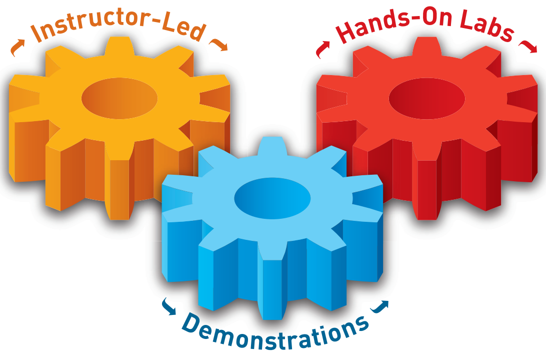

OVERVIEW

Extended Firewall Management is the next-level follow-on course to Palo Alto Networks™ Installation, Configuration, and Management (PAN-EDU-201).

Extended Firewall Management expands on 201 course topics, while introducing many new features and functions of Palo Alto Networks Next-Generation firewalls.

COURSE OBJECTIVES

Successful completion of this two-day, instructor-led course will enhance the student's understanding of how to install, configure, manage, and perform basic troubleshooting on the entire line of Palo Alto Networks Next-Generation firewalls.

Additionally, students will be instructed on the basics of implementing and managing GlobalProtect and Active/Active High Availability.

Students will gain an in-depth knowledge of how to optimize their visibility and control over applications, users, and content.

SCOPE

- **Course level:** Introductory
- **Course duration:** 2 Days
- **Course format:** Combines Instructor-led and hands-on labs
- **Platform support:** All Palo Alto Networks next-generation firewall models running PAN-OS 5.0.

TARGET AUDIENCE

- Security Engineers, Network Engineers, and Support staff

PREREQUISITES:

- Completion of Firewall Installation, Configuration, and Management (201) or equivalent experience is highly recommended
- Students must have a basic familiarity with networking concepts including routing, switching, IP addressing, and basic port-based security concepts.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

PAN_DS_EDU205_010213

Exhibit 24

ACCREDITED CONFIGURATION ENGINEER (ACE)

The Accredited Configuration Engineer (ACE) exam tests your knowledge of the core features and functions of Palo Alto Networks next-generation firewalls. The ACE exam is web-based and consists of 72 multiple-choice questions. The exam is not timed, and you can retake it as many times as necessary to earn a passing score.

EXAM OBJECTIVES

The primary goal of the ACE exam is to serve as an objective indication of your ability to configure Palo Alto Networks firewalls using the PAN-OS. To pass the ACE exam, you need a foundational knowledge of, and hands-on familiarity, with PAN-OS configuration.

REGISTRATION

Palo Alto Networks is pleased to announce its new Learning Center. Whether you're a network security engineer or a sales executive for a reseller partner, the new learning center will become a primary resource for managing your Palo Alto Networks transcript and credentials. Register Now!

RECERTIFICATION REQUIREMENTS

There are no recertification requirements for you. Currently, the ACE accreditation is tied to a specific PAN-OS version. For example, the ACE 5.0 exam is based on PAN-OS version 5.0. Partners may have recertification requirements.

BENEFITS

Passing the ACE exam indicates that you possess the basic knowledge to successfully configure Palo Alto Networks firewalls using the PAN-OS. The exam also serves as a study aid for you to prepare for CNSE certification. Finally, the ACE exam is a requirement for access to the migration tool.

FAQS

For additional details about the ACE exam, please refer to the ACE frequently asked questions (ACE FAQs).

Exhibit 25

CERTIFIED NETWORK SECURITY ENGINEER (CNSE)

cnse
certified network security engineer

The Certified Network Security Engineer (CNSE) exam is a formal certification exam that is hosted and proctored by a third-party testing company, Kryterion. The CNSE exam should be taken by anyone who wants to demonstrate a deep understanding of Palo Alto Networks technologies. This includes anyone who uses Palo Alto Networks products such as partners, system engineers, systems integrators, and support engineers.

EXAM OBJECTIVES

The Certified Network Security Engineer (CNSE) is a formal, third-party proctored certification, which indicates that those who have passed it possess an in-depth engineering level knowledge of how to install, configure, and implement Palo Alto Networks products. The CNSE exam will always be tied to the .1 release of each version of PAN-OS (4.1, 5.1, etc.). It should be taken by anyone who wishes to demonstrate a deep understanding of Palo Alto Networks technologies. This includes customers who use Palo Alto Networks products, value-added resellers, pre-sales system engineers, system integrators and varied tiers of support staff.

CONTACT

For questions or issues related to the CSNE exam please contact: cnse@paloaltonetworks.com

REGISTRATION

To take the CNSE, click here to be taken to the Kryterion registration site

TEST CENTERS

Click here to locate the test center most convenient for you.

RECERTIFICATION REQUIREMENTS

There are no recertification requirements. Currently, the CNSE certification is tied to a specific PAN-OS version. For example, the CNSE 4.1 exam is based on PAN-OS version 4.1. Partners may have recertification requirements.

STUDY GUIDE

Download a PDF with information used to prepare for the CNSE exam.

TECH DOCUMENTS

Download a zip file containing documents that support the CNSE Study Guide and test taker's exam preparation efforts.

FAQS

For additional details about the CNSE exam, please refer to the CNSE frequently asked questions (CNSE FAQs)

Exhibit 26

Consulting Services: A Part of the Solution Assurance Services Framework

Consulting Services are foundational to our approach for helping customers achieve the highest degree of protection, and remain safe from security threats as the technology landscape changes and business needs evolve. To accomplish this, we offer several consulting options to support all phases of your project lifecycle.

These include:

- **Deployment Assistance:** Deployment assistance helps organizations accelerate the time to production, reduces risk through design assistance, and full-featured implementation.
- **Management Assistance:** Management assistance mitigates the risk of disruption and improves outcomes, when configuration changes are necessary
- **Optimization Assistance:** Optimization assistance provides optimal performance and maximum protection over time

Our consultants are experienced network security professionals and Palo Alto Networks product specialists. They are supported by processes, methodologies, and best practices, which are proven successful in thousands of customer engagements. Our consulting services address the specific needs of our customers in order to deliver successful outcomes. This is why our most successful customers leverage Consulting Services by Palo Alto Networks® the network security company™.

> *"Palo Alto Networks took time to understand our business and our requirements. They went the extra mile to be there when we needed them, and helped us get our project done on time."*
>
> **Melissa Tony**
>
> **Data Communication Specialist Lead - Network Security Team**
>
> **University of Pittsburgh Medical Center**

DEPLOYMENT ASSISTANCE

Consulting services can help you with deployment assistance in order to achieve a more complete and reliable implementation to protect your business and deliver results.

Product Implementation Service

Experienced consultants from Palo Alto Networks provide on-site personalized assistance to create the optimal implementation for your business. Armed with extensive knowledge about next-generation firewall product features, and implementation best practices, our consultants will ensure your business is reliably protected from the beginning, by following these steps:

- Business and Technology Assessment
- Solution and Implementation Design
- Install and Configure
- MIGRATE (a process for replacing legacy systems with Palo Alto Networks)
- Promote to Production
- Production Audit, Tuning, and Completion

paloalto NETWORKS
the network security company™

Beginning with a detailed assessment of your technology environment (security architecture, rules, and other technology) and business requirements, we will lead the effort to design, install, configure, and tune your next-generation firewall.

With our unique MIGRATE capabilities, we can automatically leverage extensive security rules already installed in your legacy firewall systems, without having to reconfigure them.

Finally, we will support you during the testing and acceptance phases and provide expert guidance during the final move to production. Our team will schedule a post-production implementation review to confirm the project was successful, address remaining issues, and fine-tune the final configuration.

With this process, customers save weeks or even months of time that would have otherwise been spent learning, installing, and evolving the system to meet production standards.

Remote Install with Baseline Threat Protection

Installing network security products are critical events, where time, experience and expertise are essential for success. For many organizations, however, staff engineers are consumed with several projects and initiatives, and may not have the time to attend training or promptly install new systems. With time and staff in short supply, it may be challenging to dedicate the proper resources for a successful deployment.

To help customers overcome these challenges, we offer Remote Install with Baseline Threat Protection—a cost effective option to quickly (and properly) install the next-generation firewall with a standard set of options, thus providing the path to accelerate deployment.

This service includes the following:

- Pre-installation planning with customer staff
- Installation, licensing, and product registration
- Interface configuration
- Policy configuration (2 per zone)
- Knowledge transfer—overview of the installation and key product features

Using a proven methodology for planning, installation, and configuration, our experienced installation consultant will ensure that each step of the process is done promptly and accurately. After installation and configuration, the consultant will provide basic knowledge transfer to your project team members to help them manage the system and assume administrative responsibilities.

CHANGE MANAGEMENT SERVICES

Change Impact Analysis and Validation Service

Every security organization must be prepared to manage change. Change can take the form of new security requirements, new technologies, and product updates. Change, from time to time, is inevitable, and spans the range from basic configuration updates to full PAN-OS upgrades.

However, change can introduce concerns about risk. As part of a mission critical initiative, organizations need to carefully evaluate proposed changes that will perform as expected and without unintended consequences.

Change Impact Analysis and Validation Service by Palo Alto Networks helps customers fully understand the impact of planned changes, to minimize risk and ensure successful results. With this service, our next-generation firewall experts will partner with you to execute the following steps:

- Evaluate Planned Changes
- Develop a Test Plan for Change Validation
- Replicate Customer Environment and Apply Changes
- Execute Test Plan and Communicate the Results

The consultant will review the business reasons behind the proposed changes and discuss important design principles to fully understand and document the requirements. Our experienced consultants will apply product and best practices experience to establish an appropriate test plan, which will include pass/fail criteria according to your objectives. We can then quickly establish a test environment based on your current production system and settings, and even using our own lab systems when needed to accelerate the process. Once the consultant has finished executing the test plan, results will be analyzed, and a detailed report will be provided outlining our findings and next step recommendations.

Change Impact Analysis and Validation Services by Palo Alto Networks helps customers eliminate unnecessary changes, and validate the impact of proposed changes before moving into production, thus saving time and reducing exposure to risk.

OPTIMIZATION SERVICES

Successful organizations are not idle once a technology solution is operational. Both business requirements and security challenges evolve over time, and the solution must adapt to address new needs. This is especially true for network security initiatives, which operate in a constantly changing landscape of technology and threats.

Keep your next-generation firewalls operating at peak performance with Optimization Services by Palo Alto Networks. Our experienced consultants will apply product expertise and best practices knowledge to evaluate and optimize your next-generation firewall system.

Configuration Review and Optimization Service

The Configuration Review and Optimization Service involves three specific steps to uncover and address opportunities to optimize a deployment. At each step in the process, our experienced consultants will apply their extensive knowledge of Palo Alto Networks next-generation firewalls and best practices to identify recommended changes, which when implemented, will lead to optimal results.

Architecture Review: The consultant partners with the customer to review architecture and topological design of your next-generation firewall. The consultant will validate the appropriate platform for current and projected performance requirements and will help determine a future migration plan for when requirements change. In addition, the consultant review management and log monitoring options available from Palo Alto Networks and technology partners.

System Health Check: The consultant will analyze the existing production system and gather performance metrics. The consultant will review firewall system parameters, such as sessions, resources, and drops to verify that the firewall is performing optimally. Additional checks will review traffic, threat, and system logs to identify recommended changes where applicable.

Configuration Audit: Our consultant will review the existing or proposed next-generation firewall configuration. We will compare the configuration with current best practices and provide feedback if there are optimizations available to increase performance. The consultant will also validate the configuration of the next-generation firewall and suggest ways to make the best use of various features

Product Tuning and Configuration Change Implementation (optional): If needed, our consultants can incorporate all recommended changes into a separate statement of work, and then implement them for the customer. This may include configuration and policy changes, enabling improved analytics, and more.

With the Health Check and Configuration Audit service, organizations can get peace of mind that they are getting the optimal levels of protection and performance from the Palo Alto Networks next generation firewall.

About Solution Assurance Services by Palo Alto Networks

Consulting services are in important part of the Palo Alto Networks Solution Assurance Services framework for deploying and maintaining the next-generation firewall.

With this in mind, Palo Alto Networks offers Solution Assurance Services, designed to help customers get maximum protection and value out of their investment from the beginning of their projects. By providing the right combination of consulting, education, and support services, along with proprietary product and best practices expertise, Solution Assurance Services help organizations optimize every phase of firewall implementation—from pre-installation to ongoing production management.

Save Time, Reduce Risk, Maximize Value. Let us help you maximize functionality, reliability, and availability in order and achieve overall success and satisfaction in order to achieve the full potential benefit with your next-generation firewall. Deploy, manage, and optimize your network security project with Palo Alto Networks Solution Assurance Services.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

Exhibit 27

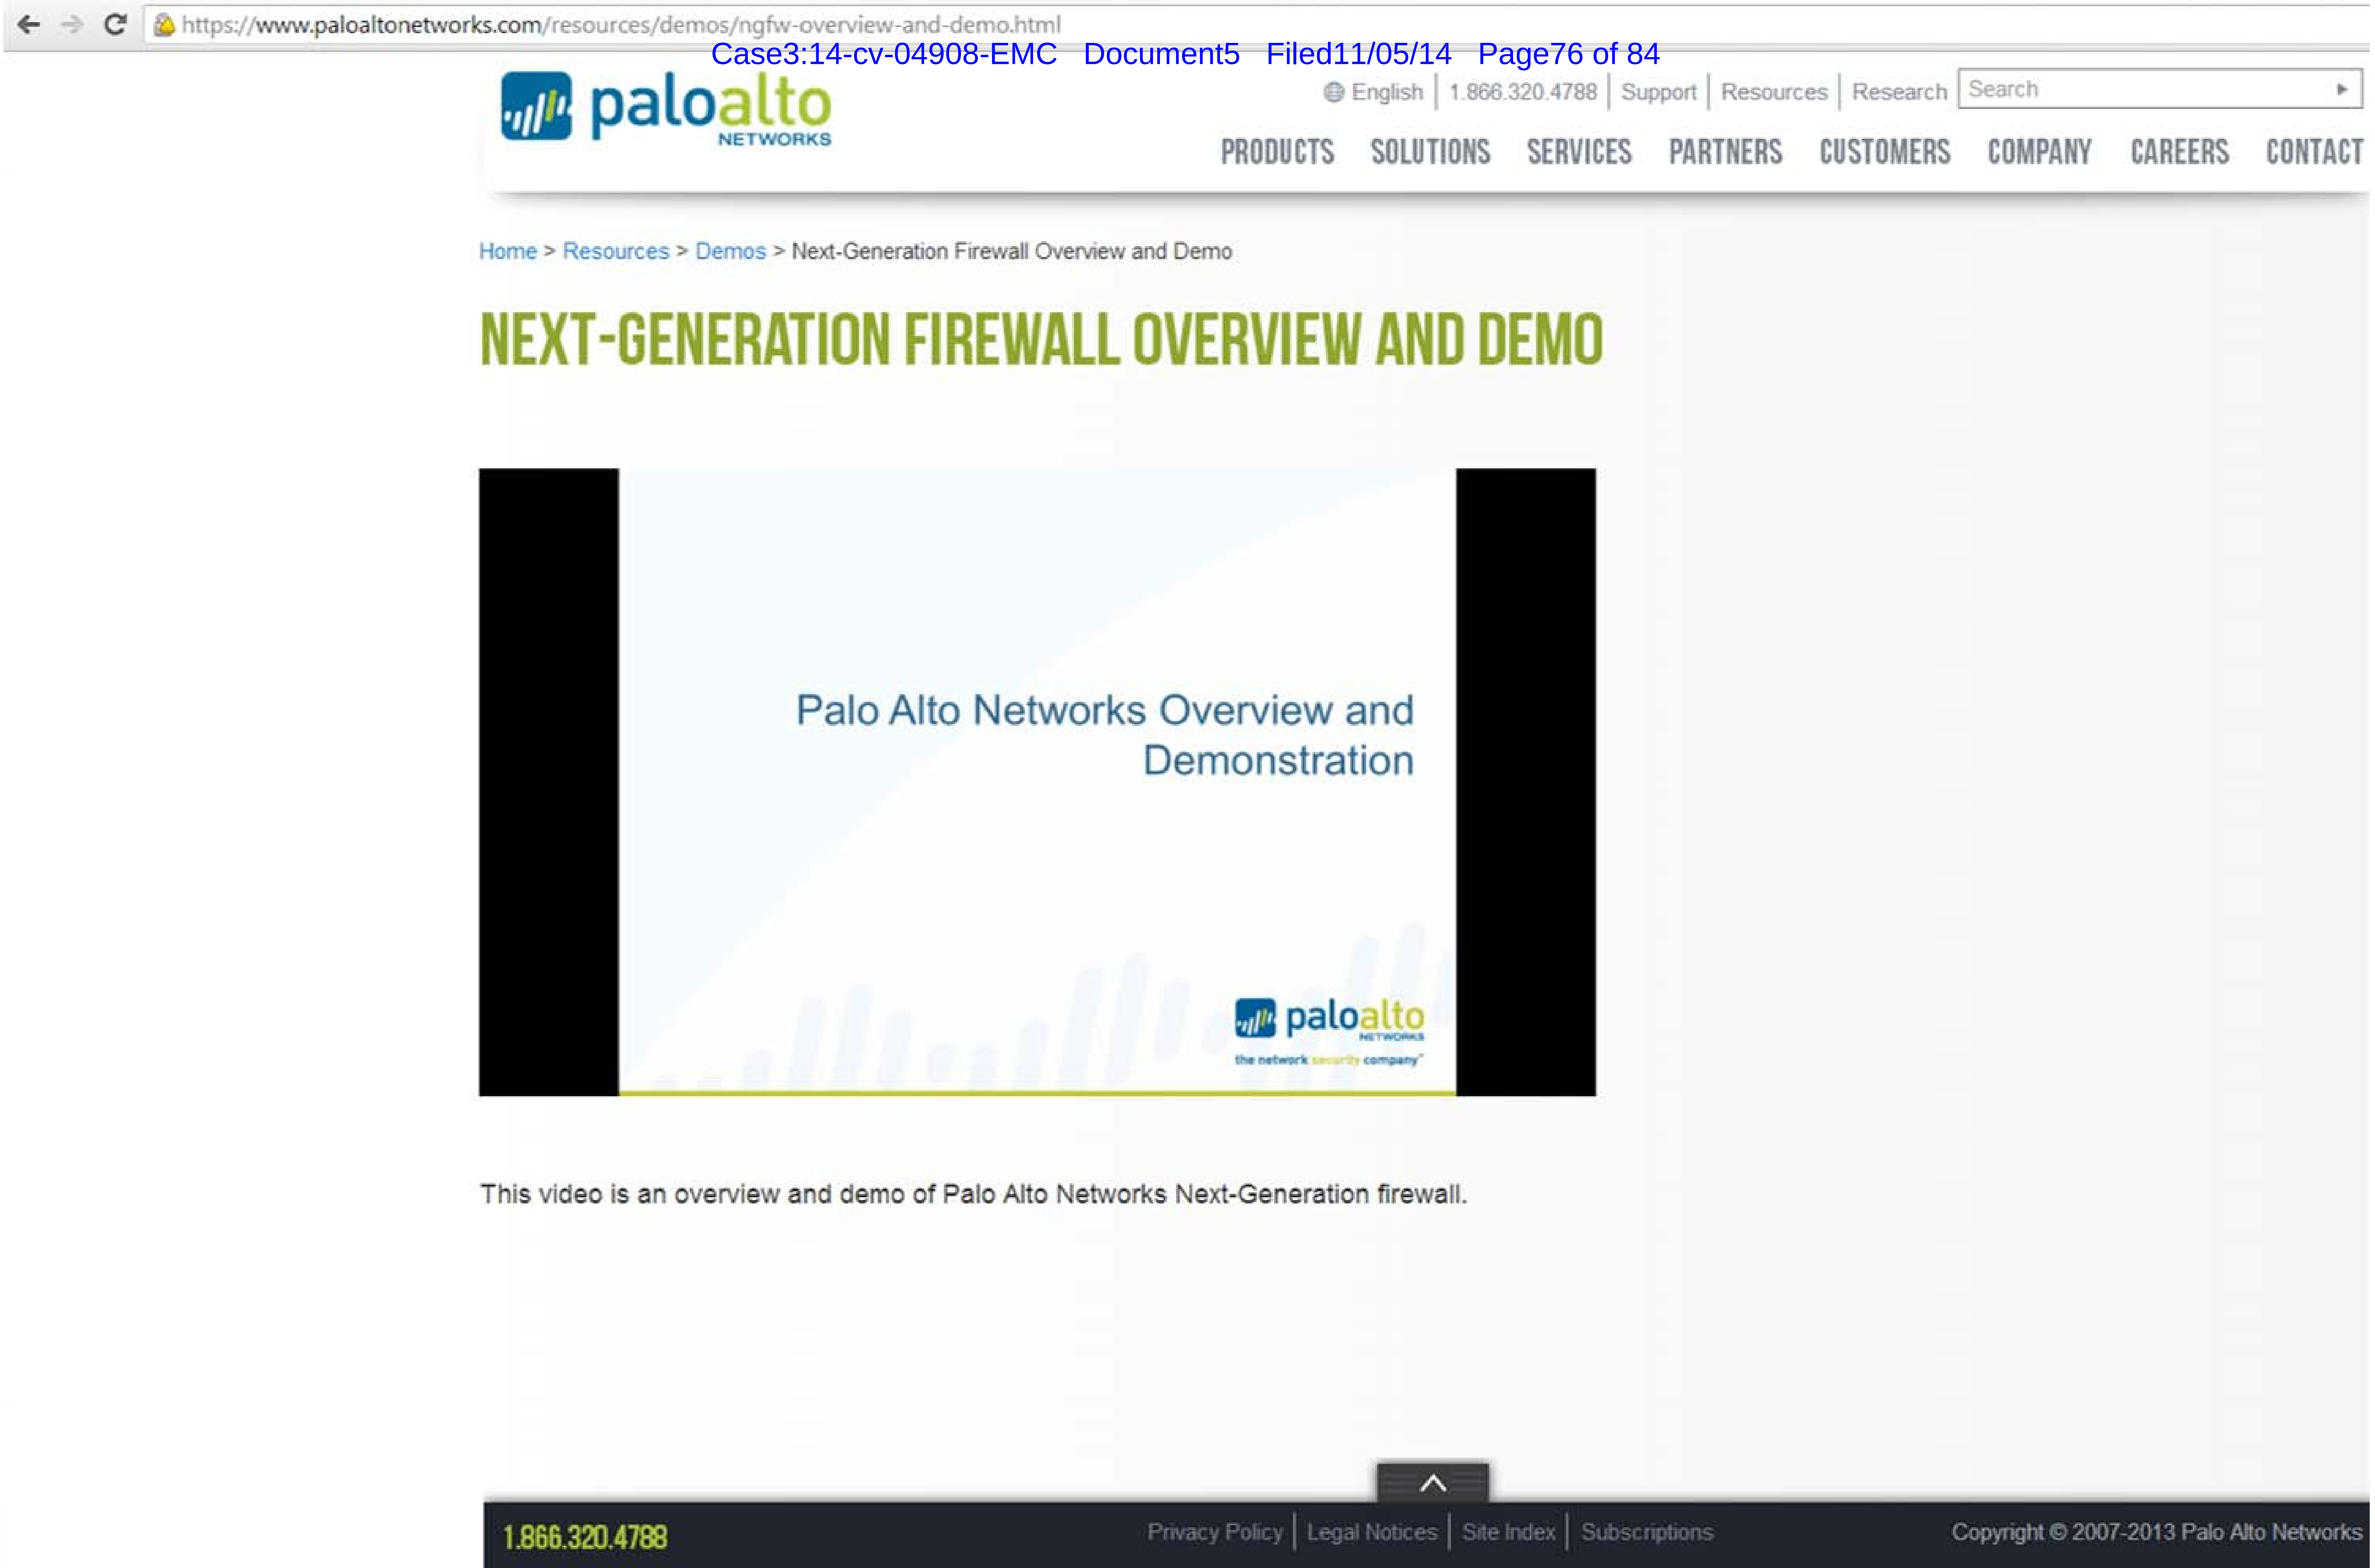

https://www.paloaltonetworks.com/resources/demos/ngfw-overview-and-demo.html

paloalto NETWORKS
English | 1.866.320.4788 | Support | Resources | Research
Search
PRODUCTS SOLUTIONS SERVICES PARTNERS CUSTOMERS COMPANY CAREERS CONTACT

Home > Resources > Demos > Next-Generation Firewall Overview and Demo

NEXT-GENERATION FIREWALL OVERVIEW AND DEMO

Palo Alto Networks Overview and Demonstration

paloalto NETWORKS
the network security company

This video is an overview and demo of Palo Alto Networks Next-Generation firewall.

1.866.320.4788
Privacy Policy | Legal Notices | Site Index | Subscriptions

Exhibit 28

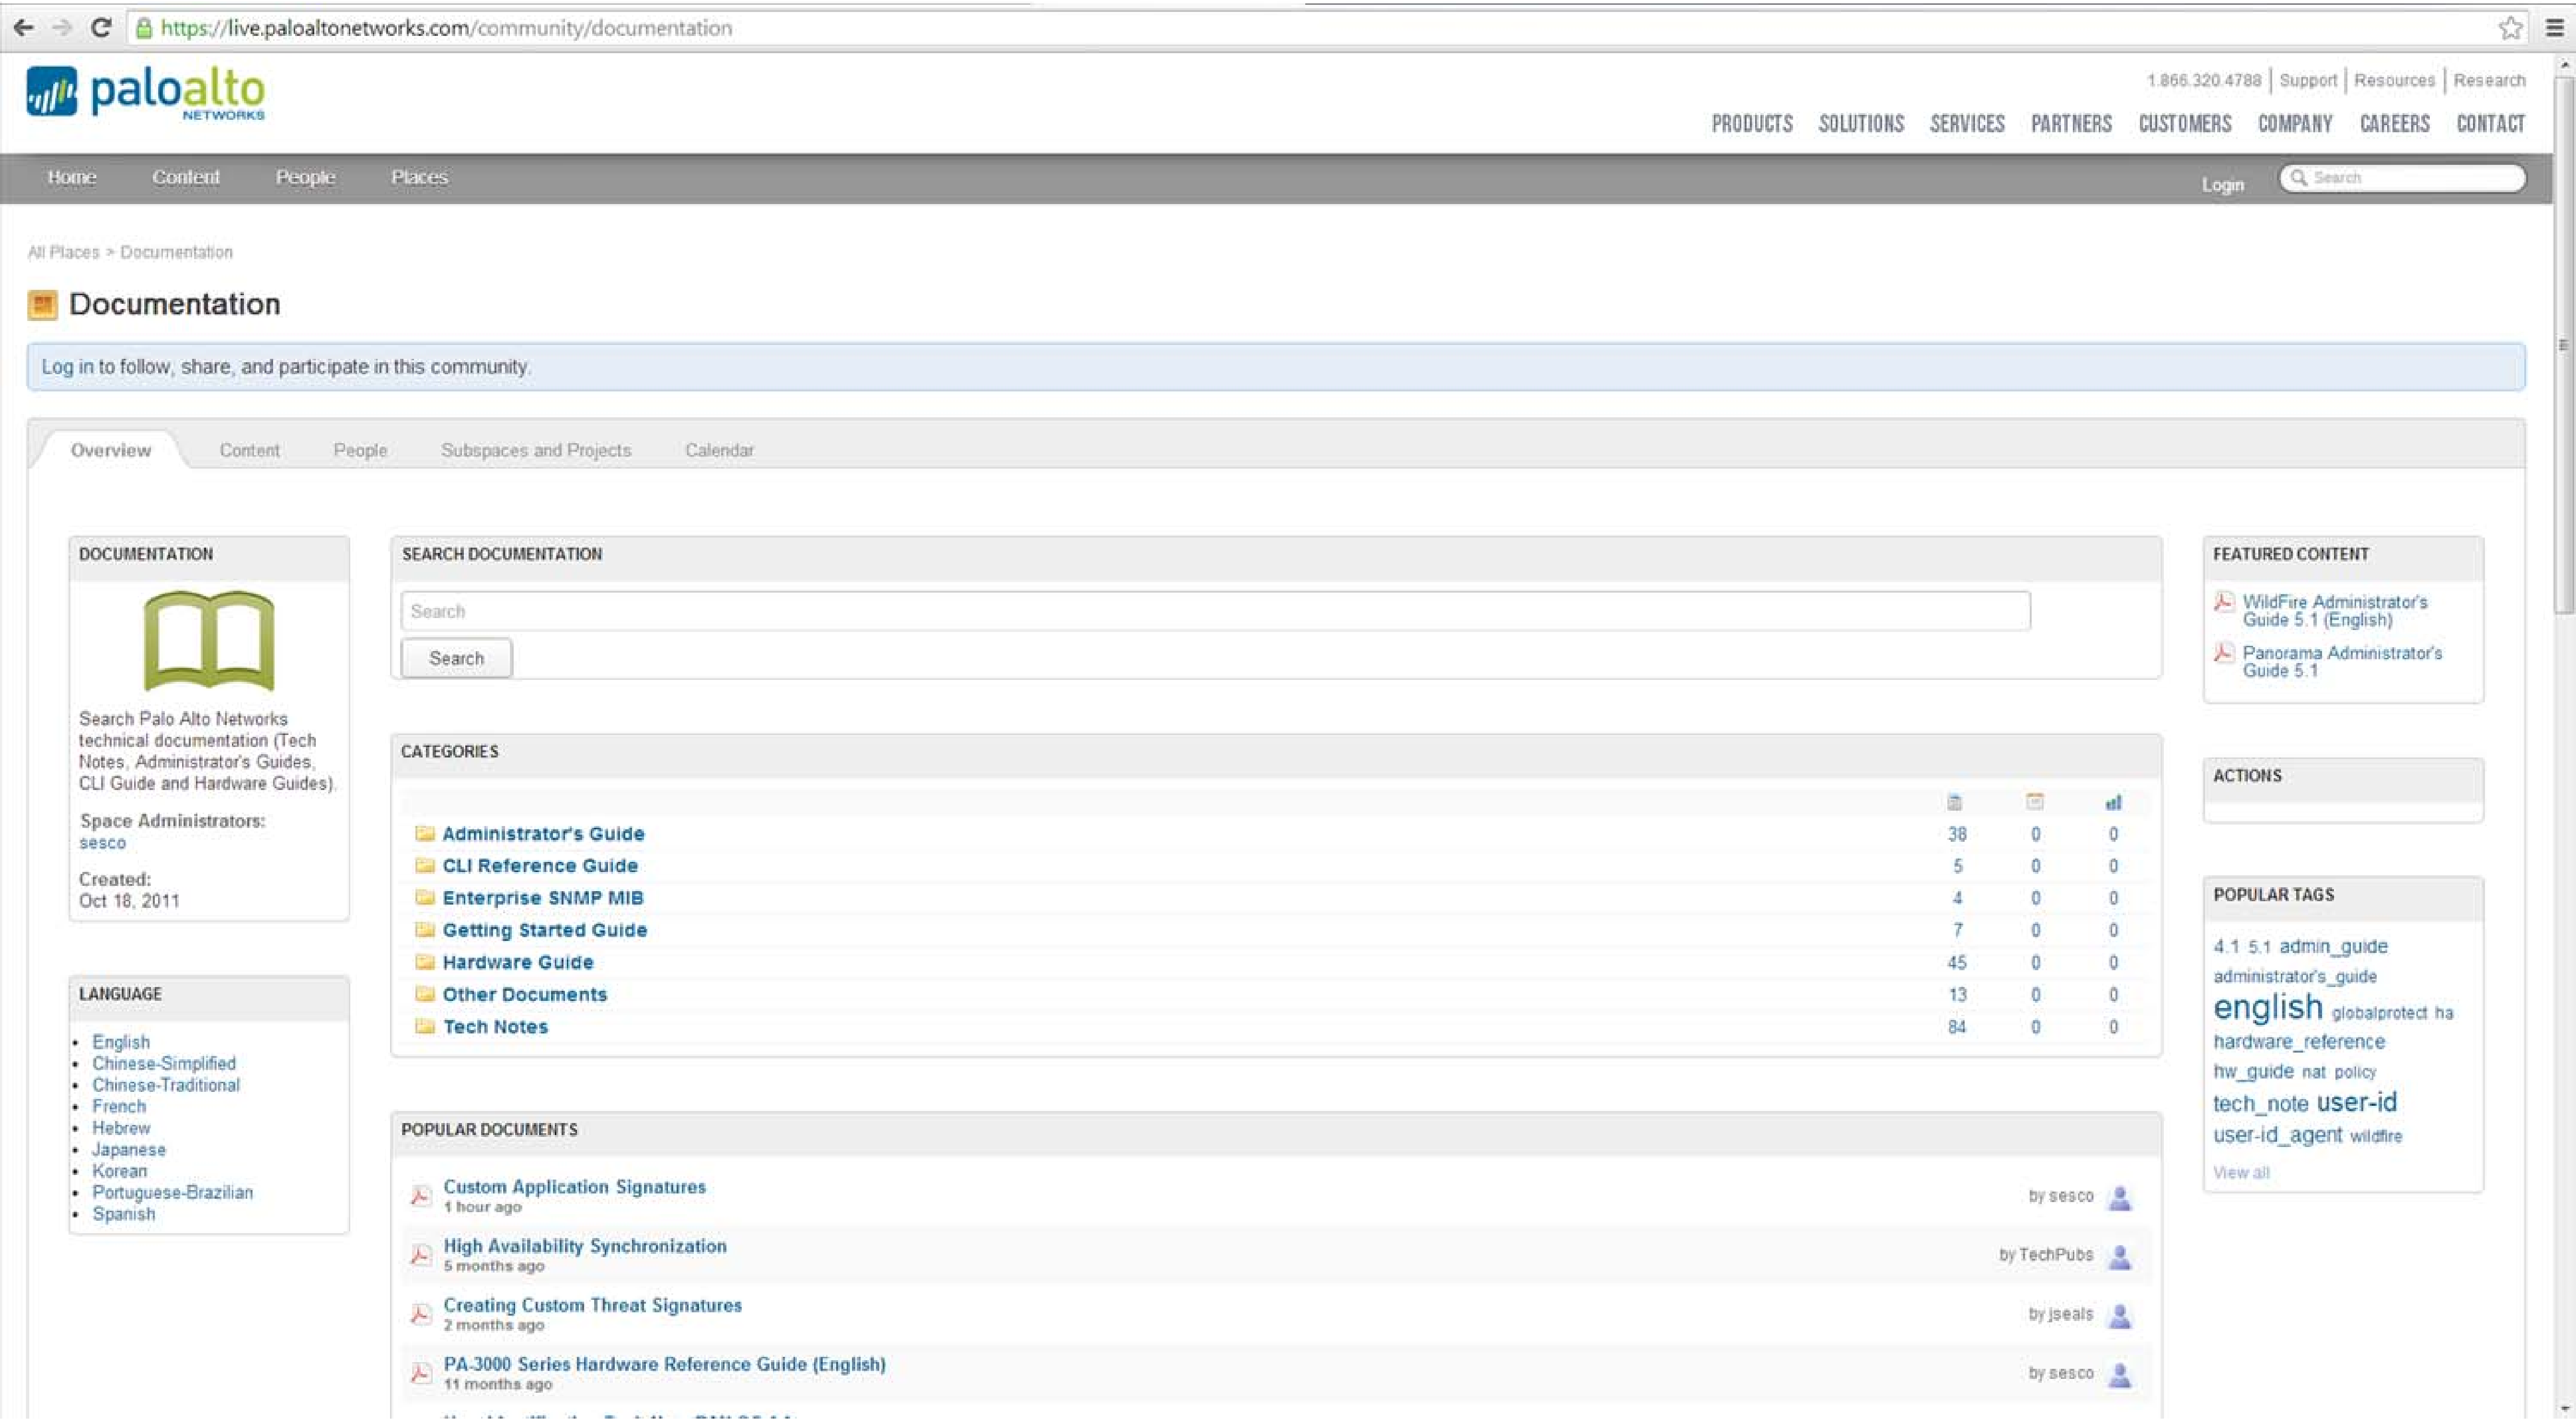
https://live.paloaltonetworks.com/community/documentation
paloalto NETWORKS
1.866.320.4788 | Support | Resources | Research
PRODUCTS SOLUTIONS SERVICES PARTNERS CUSTOMERS COMPANY CAREERS CONTACT
Home Content People Places
Login
Search

All Places > Documentation

Documentation

Log in to follow, share, and participate in this community.

Overview Content People Subspaces and Projects Calendar

DOCUMENTATION

Search Palo Alto Networks technical documentation (Tech Notes, Administrator's Guides, CLI Guide and Hardware Guides).

Space Administrators:
sesco

Created:
Oct 18, 2011

LANGUAGE

- English
- Chinese-Simplified
- Chinese-Traditional
- French
- Hebrew
- Japanese
- Korean
- Portuguese-Brazilian
- Spanish

SEARCH DOCUMENTATION

CATEGORIES

Administrator's Guide	38	0	0
CLI Reference Guide	5	0	0
Enterprise SNMP MIB	4	0	0
Getting Started Guide	7	0	0
Hardware Guide	45	0	0
Other Documents	13	0	0
Tech Notes	84	0	0

POPULAR DOCUMENTS

Custom Application Signatures
1 hour ago
by sesco

High Availability Synchronization
5 months ago
by TechPubs

Creating Custom Threat Signatures
2 months ago
by jseals

PA-3000 Series Hardware Reference Guide (English)
11 months ago
by sesco

FEATURED CONTENT

WildFire Administrator's Guide 5.1 (English)

Panorama Administrator's Guide 5.1

ACTIONS

POPULAR TAGS

4.1 5.1 admin_guide administrator's_guide english globalprotect ha hardware_reference hw_guide nat policy tech_note user-id user-id_agent wildfire

View all

Exhibit 29

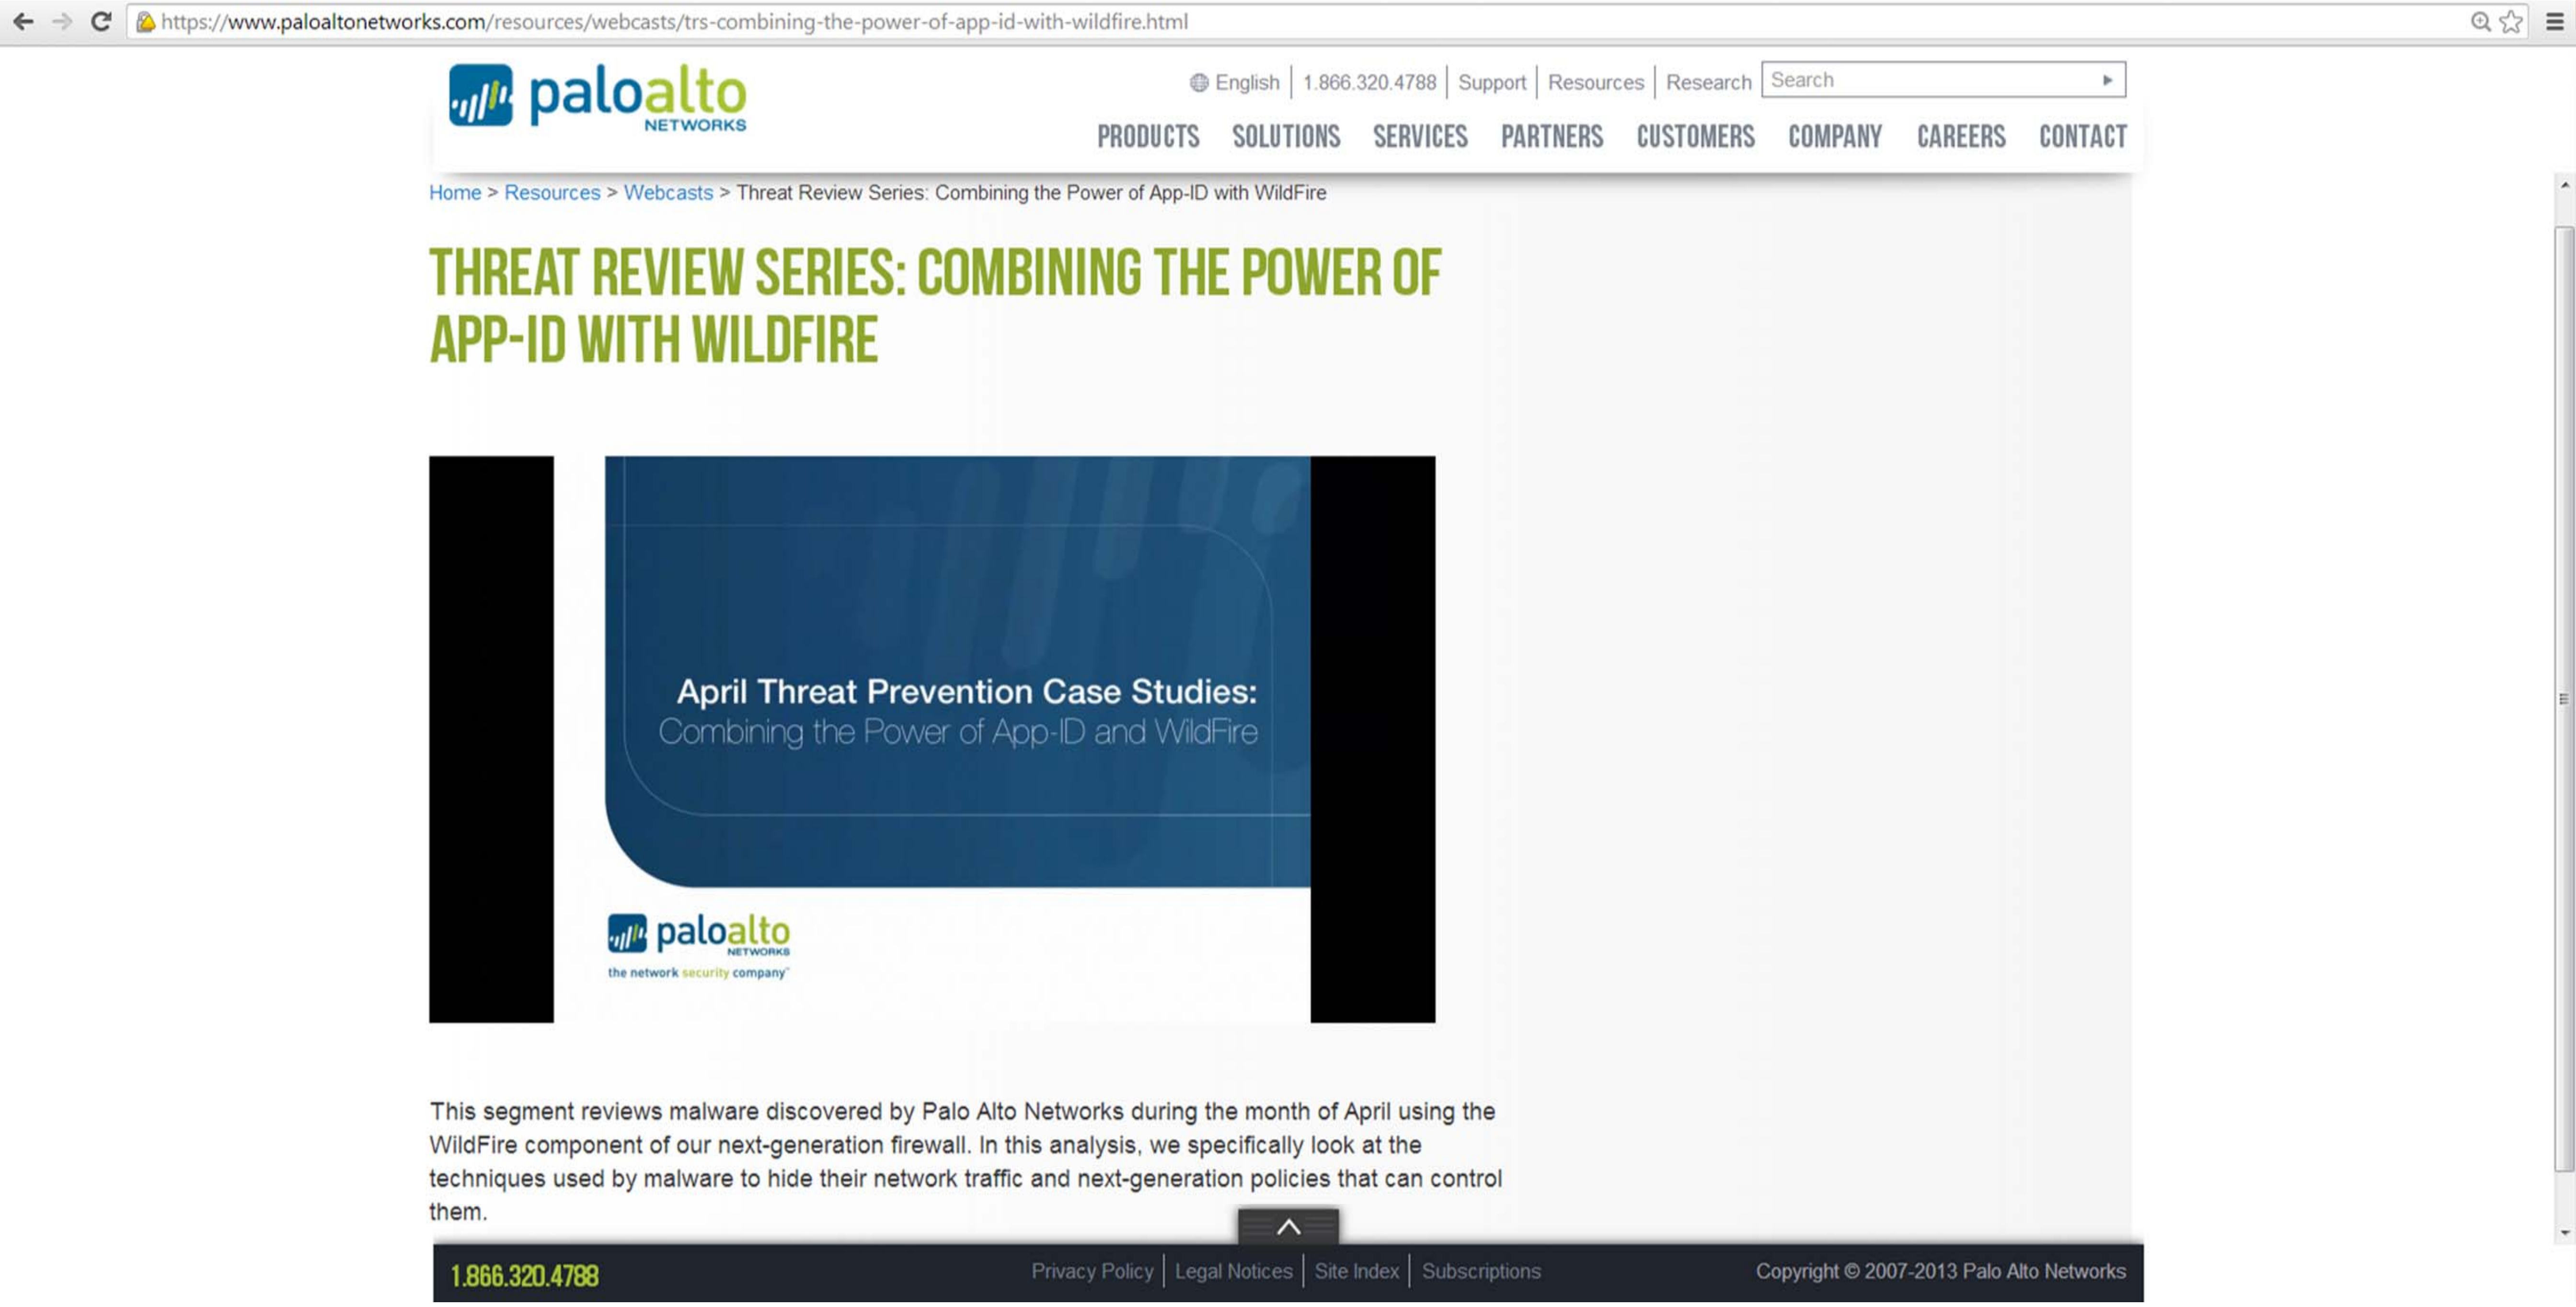

https://www.paloaltonetworks.com/resources/webcasts/trs-combining-the-power-of-app-id-with-wildfire.html

English | 1.866.320.4788 | Support | Resources | Research | Search

PRODUCTS SOLUTIONS SERVICES PARTNERS CUSTOMERS COMPANY CAREERS CONTACT

Home > Resources > Webcasts > Threat Review Series: Combining the Power of App-ID with WildFire

THREAT REVIEW SERIES: COMBINING THE POWER OF APP-ID WITH WILDFIRE

April Threat Prevention Case Studies:
Combining the Power of App-ID and WildFire

This segment reviews malware discovered by Palo Alto Networks during the month of April using the WildFire component of our next-generation firewall. In this analysis, we specifically look at the techniques used by malware to hide their network traffic and next-generation policies that can control them.

1.866.320.4788 Privacy Policy | Legal Notices | Site Index | Subscriptions Copyright © 2007-2013 Palo Alto Networks

Exhibit 30

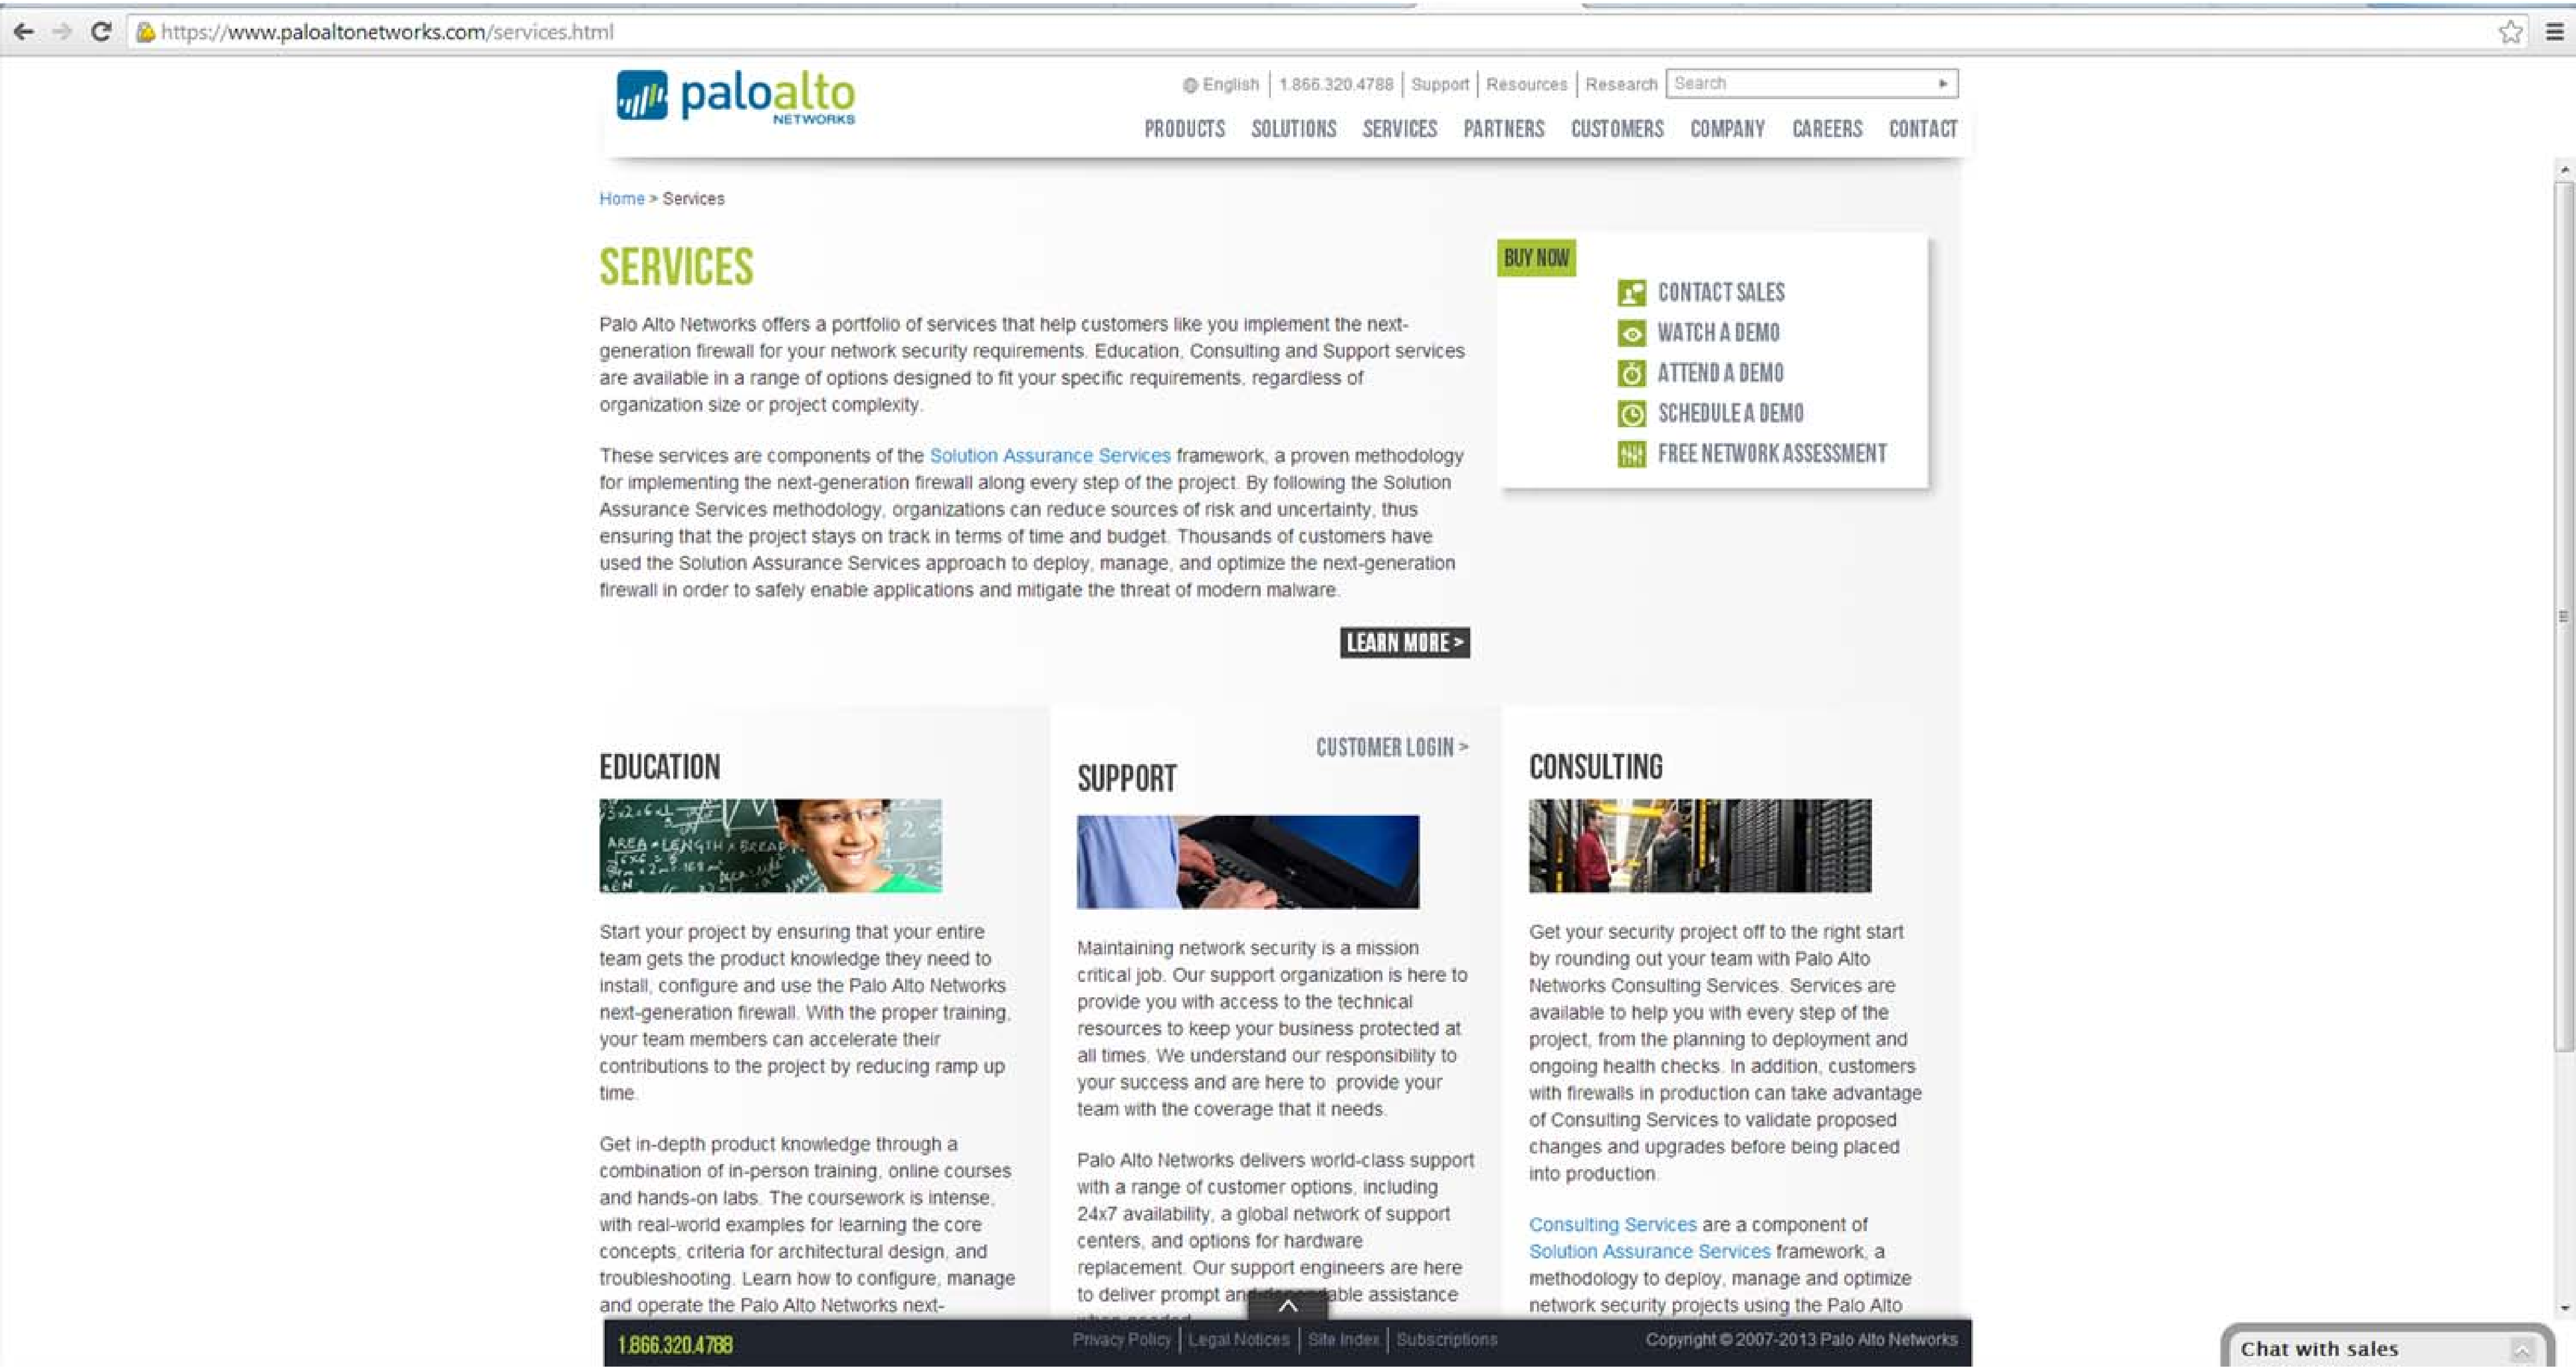

https://www.paloaltonetworks.com/services.html

paloalto NETWORKS

English | 1.866.320.4788 | Support | Resources | Research | Search

PRODUCTS SOLUTIONS SERVICES PARTNERS CUSTOMERS COMPANY CAREERS CONTACT

Home > Services

SERVICES

Palo Alto Networks offers a portfolio of services that help customers like you implement the next-generation firewall for your network security requirements. Education, Consulting and Support services are available in a range of options designed to fit your specific requirements, regardless of organization size or project complexity.

These services are components of the Solution Assurance Services framework, a proven methodology for implementing the next-generation firewall along every step of the project. By following the Solution Assurance Services methodology, organizations can reduce sources of risk and uncertainty, thus ensuring that the project stays on track in terms of time and budget. Thousands of customers have used the Solution Assurance Services approach to deploy, manage, and optimize the next-generation firewall in order to safely enable applications and mitigate the threat of modern malware.

LEARN MORE >

BUY NOW

- CONTACT SALES
- WATCH A DEMO
- ATTEND A DEMO
- SCHEDULE A DEMO
- FREE NETWORK ASSESSMENT

EDUCATION

Start your project by ensuring that your entire team gets the product knowledge they need to install, configure and use the Palo Alto Networks next-generation firewall. With the proper training, your team members can accelerate their contributions to the project by reducing ramp up time.

Get in-depth product knowledge through a combination of in-person training, online courses and hands-on labs. The coursework is intense, with real-world examples for learning the core concepts, criteria for architectural design, and troubleshooting. Learn how to configure, manage and operate the Palo Alto Networks next-

SUPPORT

CUSTOMER LOGIN >

Maintaining network security is a mission critical job. Our support organization is here to provide you with access to the technical resources to keep your business protected at all times. We understand our responsibility to your success and are here to provide your team with the coverage that it needs.

Palo Alto Networks delivers world-class support with a range of customer options, including 24x7 availability, a global network of support centers, and options for hardware replacement. Our support engineers are here to deliver prompt an[illegible]able assistance when needed.

CONSULTING

Get your security project off to the right start by rounding out your team with Palo Alto Networks Consulting Services. Services are available to help you with every step of the project, from the planning to deployment and ongoing health checks. In addition, customers with firewalls in production can take advantage of Consulting Services to validate proposed changes and upgrades before being placed into production.

Consulting Services are a component of Solution Assurance Services framework, a methodology to deploy, manage and optimize network security projects using the Palo Alto

1.866.320.4788 Privacy Policy | Legal Notices | Site Index | Subscriptions Copyright © 2007-2013 Palo Alto Networks

Chat with sales

Exhibit 31

paloalto networks.

Home > Customers >

Customers

You could take our word that our next-generation firewalls deliver unmatched visibility and control to safely enable your applications, but we think the words of our customers are more convincing. Read what they have to say about real-world deployments of our products. Below you will find case studies detailing our next-generation firewalls in action at companies large and small - in simple to complex global networks - in a variety of industries.

Filter Results By

Filter by Type

- Customer Story (99)
- Customer Video (12)
- Video (9)

Filter by Topic

- Network Security (4)
- Cybersecurity (2)
- Advanced Endpoint Protection (1)
- Endpoint Security (1)
- Firewalls (1)
- Threat (1)

Filter by Industry

- Education (27)
- Professional Services (14)
- Healthcare (9)
- Financial Services (7)
- Manufacturing (7)
- Government (6)
- High Tech (6)
- Non-Profit (6)
- Service Provider & Telecommunications (6)
- Media & Entertainment (5)

Viewing All Results

Most Viewed

CAME Group Saves US $2.5 Million in Costs and Eliminates Over 100 Devices From Network by Centralizing on Palo Alto Networks® Next-generation Security Platforms

CAMEGROUP

CAME Group, June 19, 2014, 10:00 AM

CAME Group operates in 118 countries through 480 branches and licensed dealers. Thanks to its Bpt and Urbaco brands, it is a key global player in the home automation, urban planning, and high-security sectors, for which it offers integrated solutions for regulating and monitoring people flows and access points. CAME Group does 70 percent of its business globally. It is extremely proud of its Italian heritage, and employs 1,200 staff with sales around 215 million euros in 2013. This Case Study available in: Italian.

Share 11 391 Views

1-25 of 99 ◀ 1 2 3 4 ▶

⊗ Customer Story

Sales: 866.320.4788

Privacy Policy | Legal Notices | Site Index | Subscriptions |